

# **Tri-Laboratory Linux Capacity Cluster 2 (TLCC2)**

## **Draft Statement of Work B590550 Attachment 2**

---

**Version 2.4  
August 19, 2010**



## TLCC2 Statement of Work

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor Lawrence Livermore National Security, LLC (LLNS) nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or LLNS. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or LLNS, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Security, LLC under Prime Contract No. DE-AC52-07NA27344.

# Contents

## Table of Contents

<b>1</b>	<b>Background.....</b>	<b>7</b>
1.1	Advanced Simulation and Computing Program.....	7
1.2	ASC Capacity Systems Strategy.....	7
1.3	Tri-Laboratory Simulation Environments .....	8
1.3.1	<i>The LLNL Simulation Environment</i> .....	8
1.3.2	<i>The LANL Simulation Environment</i> .....	9
1.3.3	<i>The SNL Simulation Environment</i> .....	10
1.4	Utilization of Existing Facilities.....	11
1.4.1	<i>The LLNL Existing Facilities</i> .....	12
1.4.2	<i>The LANL Existing Facilities</i> .....	15
1.4.3	<i>The Sandia Existing Facilities</i> .....	17
<b>2</b>	<b>TLCC2 Scalable Unit Strategy and Architecture .....</b>	<b>19</b>
2.1	TLCC2 Strategy .....	19
2.2	TLCC2 Cluster Architecture.....	20
2.3	TLCC2 Software Environment .....	24
2.3.1	<i>TLCC2 Synthetic Workload Test Plan</i> .....	25
2.3.2	<i>Description of Common Computing Environment (CCE)</i> .....	26
2.3.2.1	Tri-Laboratory Operating System Stack (TOSS) .....	27
2.3.2.2	Simple Linux Utility for Resource Management.....	30
2.3.3	<i>Moab Scheduler</i> .....	31
2.3.4	<i>Lustre Cluster Wide File System (Sandia, LLNL)</i> .....	31
2.3.5	<i>Panasas PanFS Multi-Cluster Environment Wide Global Parallel OBSD Based File System (LANL)</i> 31	
2.3.6	<i>Production Compiler Suites</i> .....	32
2.3.7	<i>TotalView Debugger</i> .....	32
<b>3</b>	<b>TLCC2 Technical Requirements .....</b>	<b>33</b>
3.1	High-Level Hardware Summary (TR-1).....	34
3.1.1	<i>SU High-Level Architecture</i> .....	34
3.1.2	<i>SU Requirements Summary Matrix</i> .....	35
3.1.3	<i>SU Evolution Overview (TR-1)</i> .....	37
3.2	SU Hardware Requirements.....	38
3.2.1	<i>TLCC2 Scalable Unit (MR)</i> .....	38
3.2.2	<i>TLCC2 SMP two Socket Configuration (TR-1)</i> .....	38
3.2.3	<i>SU Peak (TR-1)</i> .....	38
3.2.4	<i>Number of SU (MR)</i> .....	38
3.2.5	<i>TLCC2 Node Requirements</i> .....	39
3.2.5.1	Processor and Cache (TR-1).....	39
3.2.5.2	Node Floating Point Performance (TR-1).....	39
3.2.5.3	Chipset and Memory Interface (TR-2).....	39
3.2.5.4	Node Delivered Memory Performance (TR-1) .....	39
3.2.5.5	Node I/O Configuration (TR-2).....	39
3.2.5.6	Node Memory (TR-1).....	39
3.2.5.7	Node Power (TR-2) .....	40
3.2.5.8	Linux Access to Memory Error Detection (TR-1) .....	40
3.2.5.9	No Local Hard Disk (TR-1).....	40

## TLCC2 Statement of Work

3.2.5.10	Node Form Factor (TR-2) .....	40
3.2.5.11	Integrated Management Ethernet (TR-1) .....	40
3.2.5.12	Node BIOS (MR) .....	40
3.2.5.12.1	Node BIOS Type Options (MR).....	41
3.2.5.12.2	Remote Network Boot (MR).....	41
3.2.5.12.3	Remote IBA Network Boot (TR-1) .....	41
3.2.5.12.4	Node Initialization (TR-1) .....	41
3.2.5.12.5	Error/Interrupt Handling (MR).....	41
3.2.5.12.6	BIOS Security (MR).....	41
3.2.5.12.7	Plans and Process for Needed BIOS Updates (MR).....	42
3.2.5.12.8	Failsafe/Fallback Boot Image (TR-1) .....	42
3.2.5.12.9	Failsafe CMOS Parameters (TR-1).....	42
3.2.5.12.10	Serial Console after Failsafe/Fallback Booting (TR-1).....	42
3.2.5.12.11	BIOS Upgrade and Restore (TR-1).....	42
3.2.5.12.12	CMOS Parameter Manipulation (TR-1) .....	42
3.2.5.12.13	BIOS Command Line Interface (TR-2).....	43
3.2.5.12.14	Serial Console over Ethernet Support (TR-2).....	43
3.2.5.12.15	Power-On Self Test (TR-2).....	43
3.2.5.12.16	BIOS Security Verification (MR).....	43
3.2.5.13	Programmable LED(s) (TR-3).....	43
3.2.6	<i>IBA Interconnect (MR)</i> .....	43
3.2.6.1	Node HCA Functionality (TR-1) .....	43
3.2.6.2	Node Bandwidth, Latency and Throughput (TR-1).....	43
3.2.6.3	Fully Functional IBA Topology (TR-1) .....	44
3.2.6.4	IBA Cabling Pattern (TR-1) .....	44
3.2.6.5	IBA Interconnect Reliability (TR-1).....	44
3.2.6.6	Multi-SU Spine Switches (TR-1).....	45
3.2.7	<i>Remote Manageability (TR-1)</i> .....	45
3.2.7.1	Traditional Remote Management Solution (TR-1) .....	45
3.2.7.1.1	Remote Console (TR-3) .....	46
3.2.7.1.2	Remote Node Power Control over Management LAN (TR-3) .....	46
3.2.7.2	IPMI and BMC Remote Management Solution (TR-1) .....	46
3.2.7.2.1	ConMan Access to Console via Serial over LAN (TR-1) .....	47
3.2.7.2.2	LAN PowerMan Access (TR-1).....	47
3.2.7.2.3	LAN Management Access (TR-1).....	47
3.2.7.2.4	Traditional Remote Management Backup Plan (TR-2).....	47
3.2.7.2.5	Additional IPMI Security Requirements (TR-1).....	47
3.2.7.2.5.1	Bad Password Threshold (TR-1).....	47
3.2.7.2.5.2	Bad Password Monitoring (TR-1).....	48
3.2.7.3	Remote Management Solution Requirements (TR-1).....	48
3.2.7.3.1	Serial Console Redirection (TR-1) .....	48
3.2.7.3.2	Dedicated Serial Console Communications (TR-2).....	48
3.2.7.3.3	Serial Console Efficiency (TR-1).....	48
3.2.7.3.4	Flow Control (TR-1).....	48
3.2.7.3.5	Peripheral Device Firmware (TR-2) .....	48
3.2.7.3.6	Remote Network Boot Mechanism (TR-1).....	48
3.2.7.3.7	Serial Break (TR-1) .....	49
3.2.7.4	Remote Management Security Requirements (TR-1).....	49
3.2.7.5	GPU Node Requirements (MOR).....	49
3.2.7.6	GPU Node General Requirement (MOR) .....	49
3.2.7.7	GPU Node Architecture (MOR).....	49
3.2.7.8	GPU Node Memory Requirement (MOR).....	49
3.2.8	<i>Gateway Node Requirements (TR-1)</i> .....	49
3.2.8.1	Gateway Node Count (TR-1) .....	50
3.2.8.2	Gateway Node Configuration (TR-1).....	50
3.2.8.3	Gateway Node I/O Configuration (TR-1) .....	50
3.2.8.4	Gateway Node QDR IB Card (TR-1).....	50

## TLCC2 Statement of Work

3.2.8.5	Gateway Node 10Gb Ethernet Card (TR-1).....	50
3.2.8.6	Gateway Node Delivered Performance (TR-2).....	50
3.2.9	<i>Login/Service/Master Node Requirements</i> .....	51
3.2.9.1	LSM Node Count (TR-1).....	51
3.2.9.2	LSM Node I/O Configuration (TR-1) .....	51
3.2.9.3	LSM Node Ethernet Configuration (TR-1) .....	51
3.2.9.4	LSM Node Accessory Configuration (TR-2).....	51
3.2.10	<i>Remote Partition Service Node Requirements</i> .....	51
3.2.10.1	RPS Node Count (TR-1).....	51
3.2.10.2	RPS Node I/O Configuration (TR-1).....	51
3.2.10.3	RPS Node Ethernet Configuration (TR-1).....	52
3.2.10.4	RPS Node RAID Configuration (TR-1) .....	52
3.2.11	<i>SU Management Ethernet (TR-1)</i> .....	52
3.2.12	<i>SU Racks and Packaging (TR-1)</i> .....	52
3.2.12.1	SU Design Optimization (TR-2) .....	53
3.2.12.2	Rack Height and Weight (TR-1).....	53
3.2.12.3	Rack Structural Integrity (TR-2).....	53
3.2.12.4	Rack Air Flow and Cooling (TR-1) .....	53
3.2.12.5	Rack Doors (TR-2) .....	53
3.2.12.6	Rack Cable Management (TR-2) .....	54
3.2.12.7	Rack Color (TR-3).....	54
3.2.12.8	Rack Power and Cooling (TR-1).....	54
3.2.12.9	Rack PDU (TR-1).....	54
3.2.13	<i>Safety Standards and Testing (TR-1)</i> .....	55
3.3	TLCC2 Software Requirements (TR-1).....	55
3.3.1	<i>Minimum IBA Software Stack (MR)</i> .....	55
3.3.1.1	IBA Software Stack Compatibility (MR).....	55
3.3.1.2	Open Source IBA Software Stack (TR-1).....	56
3.3.2	<i>IBA Upper Layer Protocols (TR-2)</i> .....	56
3.3.2.1	TLCC2 IB HCA Error Reporting (TR-3).....	56
3.3.2.2	TLCC2 IB Switch Firmware Update (TR-3).....	56
3.3.3	<i>TLCC2 Peripheral Device Drivers (TR-1)</i> .....	56
3.3.4	<i>GPU Node Software (MOR)</i> .....	56
3.3.5	<i>RPS Node Software (TR-1)</i> .....	57
3.3.6	<i>Hardware Memory Uncorrectable Error Detection (TR-1)</i> .....	57
3.3.7	<i>Hardware Memory Corrected Error Detection (TR-1)</i> .....	57
3.3.8	<i>Hardware Memory Controller Capabilities &amp; Configuration (TR-1)</i> .....	58
3.3.9	<i>Software Support for Memory Error Detection and Configuration (TR-1)</i> .....	58
3.3.10	<i>Memory Diagnostics (TR-1)</i> .....	59
3.3.11	<i>Linux Access to Motherboard Sensors (TR-1)</i> .....	59
3.3.12	<i>Remote Management Software (TR-2)</i> .....	60
3.3.12.1	TRMS Software (TR-1) .....	60
3.3.12.2	IPMI and BMC Remote Management Software (TR-1).....	60
3.3.12.3	Linux Tool for BIOS Upgrade (TR-1) .....	60
3.3.13	<i>System Diagnostics (TR-2)</i> .....	61
3.4	SU Hardware Evolution (TR-1).....	61
3.5	SU Software Evolution (TR-1) .....	61
4	<b>Reliability, Availability, Serviceability (RAS) and Maintenance</b> .....	62
4.1	Highly Reliable Management Network (TR-1).....	63
4.2	TLCC2 Node Reliability and Monitoring (TR-1).....	63
4.3	In Place Node Service (TR-1) .....	63
4.4	Component Labeling (TR-1).....	64
4.5	Field Replaceable Unit (FRU) Diagnostics (TR-2).....	64
4.5.1	<i>Node Diagnostics Suite (TR-1)</i> .....	64

## TLCC2 Statement of Work

4.5.2	<i>Memory Diagnostics (TR-1)</i> .....	64
4.5.3	<i>IBA Diagnostics (TR-1)</i> .....	64
4.5.4	<i>IPMI Based diagnostics (TR-1)</i> .....	65
4.5.5	<i>Peripheral Component Diagnostics (TR-2)</i> .....	65
4.6	Scalable System Monitoring (TR-2).....	65
4.7	Hardware Maintenance (TR-1).....	65
4.7.1	<i>On-site Parts Cache (TR-1)</i> .....	66
4.7.1	<i>Hot Spare Cluster (TR-1)</i> .....	66
4.7.2	<i>Statement of Volatility (MR)</i> .....	67
4.8	Software Support (TR-1).....	67
4.9	Mean Time Between Failure (MTBF) Calculation (TR-1).....	67
<b>5</b>	<b>Facilities Information</b> .....	<b>68</b>
5.1	LLNL Facilities Information.....	68
5.2	LANL Facilities Information.....	68
5.3	SNL Facilities Information.....	69
5.4	Power Requirements (TR-1).....	70
5.5	Cooling Requirements (TR-1).....	70
5.6	Floor Space Requirements (TR-1).....	70
5.7	Delivery Requirements (TR-1).....	70
5.7.1	<i>SU Installation Time (TR-1)</i> .....	71
<b>6</b>	<b>Project Management</b> .....	<b>72</b>
6.1	Risk Reduction Plan (TR-1).....	72
6.2	Open Source Development Partnership (TR-2).....	72
6.3	Project Manager (TR-1).....	72
6.4	Project Milestones (TR-3).....	73
6.4.1	<i>Detailed Project Plan (TR-1)</i> .....	73
6.4.2	<i>Tri-Laboratory TOSS Final Checkout (TR-1, April 15, 2011)</i> .....	74
6.4.3	<i>TLCC2 Phase 1 Build (TR-1, June 2011)</i> .....	74
6.4.4	<i>TLCC2 SU Phase 1 Delivery and Acceptance (TR-1, July 2011)</i> .....	75
6.4.5	<i>TLCC2 Phase 1 Cluster Integration (TR-1, July 2011)</i> .....	75
6.4.6	<i>TLCC2 Phase 1 Option Build (TR-1, July 2011)</i> .....	76
6.4.7	<i>TLCC2 SU Phase 1 Option Delivery and Acceptance (TR-1, August 2011)</i> .....	76
6.4.8	<i>TLCC2 Phase 1 Option Cluster Integration (TR-1, August 2011)</i> .....	77
6.4.9	<i>TLCC2 Phase 2 SU Build (TR-1, August 2011)</i> .....	78
6.4.10	<i>TLCC2 Phase 2 SU Delivery and Acceptance (TR-1, September 2011)</i> .....	78
6.4.11	<i>TLCC2 Phase 2 Cluster Integration (TR-1, September 2011)</i> .....	79
6.4.12	<i>TLCC2 Phase 3 Build (TR-1, October 2011)</i> .....	79
6.4.13	<i>TLCC2 SU Phase 3 Delivery and Acceptance (TR-1, November 2011)</i> .....	80
6.4.14	<i>TLCC2 Phase 3 Cluster Integration (TR-1, November 2011)</i> .....	80
6.4.15	<i>TLCC2 Phase 3 Option Build (TR-1, November 2011)</i> .....	81
6.4.16	<i>TLCC2 SU Phase 3 Option Delivery and Acceptance (TR-1, December 2011)</i> .....	82
6.4.17	<i>TLCC2 Phase 3 Option Cluster Integration (TR-1, December 2011)</i> .....	82
<b>7</b>	<b>Glossary</b> .....	<b>84</b>
7.1	General.....	84
7.2	Hardware.....	84
7.3	Software.....	87

# 1 Background

## 1.1 Advanced Simulation and Computing Program

The Advanced Simulation and Computing (ASC) Program (formerly known as the Accelerated Strategic Computing Initiative, ASCI) has led the world in capability computing for the last ten years. Capability computing is defined as a world-class platform (in the Top10 of the Top500.org list) with scientific simulations running at scale on the platform. Example systems are ASCI Red, Blue-Pacific, Blue-Mountain, White, Q, RedStorm, Purple, and Roadrunner (see [http://asc.llnl.gov/computing\\_resources/](http://asc.llnl.gov/computing_resources/), <http://www.sandia.gov/NNSA/ASC/platforms.html> and <http://www.lanl.gov/roadrunner/>). ASC applications have scaled to multiple thousands of CPUs and accomplished a long list of mission milestones on these ASC capability platforms. However, the computing demands of the ASC and Stockpile Stewardship Programs also include a vast number of smaller scale runs for day-to-day simulations. Indeed, every “hero” capability run requires many hundreds to thousands of much smaller runs in preparation and post processing activities. In addition, there are many aspects of the Stockpile Stewardship Program (SSP) that can be directly accomplished with these so-called “capacity” calculations. The need for capacity is now so great within the program that it is increasingly difficult to allocate the computer resources required by the larger capability runs. To rectify the current “capacity” computing resource shortfall, the ASC program has allocated a large portion of the overall ASC platforms budget to “capacity” systems. In addition, within the next five to ten years the Life Extension Programs (LEPs) for major nuclear weapons systems must be accomplished. These LEPs and other SSP programmatic elements will further drive the need for capacity calculations and hence “capacity” systems as well as future ASC capability calculations on “capability” systems.

To respond to this workload analysis, the ASC Program is making a large sustained strategic investment in these capacity systems, which started in Government Fiscal Year 2007 (GFY07). This second Tri-Laboratory Linux Capacity Cluster (TLCC2) procurement represents a continuation the ASC Program’s investment vehicle in these capacity systems. It also builds on the previous strategy for quickly building, fielding and integrating many Linux clusters of various sizes into classified and unclassified production service through a concept of Scalable Units (SU). The programmatic objective is to dramatically reduce the overall Total Cost of Ownership (TCO) of these “capacity” systems relative to the best practices in Linux Cluster deployments today. This objective only makes sense in the context of these systems quickly becoming very robust and useful production clusters under the crushing load that will be inflicted on them by the ASC Program and SSP scientific simulation capacity workload.

## 1.2 ASC Capacity Systems Strategy

The ASC Program “capacity” systems strategy leverages the extensive experience fielding world class Linux clusters within the Tri-Laboratory community, which consists of the Los Alamos National Laboratory (LANL), the Sandia National Laboratories (SNL), and the Lawrence Livermore National Laboratory (LLNL). This strategy is based on the observation that the Commercial, off-the-shelf (COTS) marketplace is demand driven by volume purchases. As such, the TLCC procurement is designed to maximize the purchasing power of the ASC Program by changing the past Linux cluster purchasing practices. If the Tri-

Laboratories extended existing practices then each laboratory would separately procure Linux clusters (possibly multiple times per year) over multiple years. In addition, ASC Program applications developers and end-users have enjoyed a reduction in the number of different combinations of instruction set architectures, interconnects, compilers, and operating systems (OS's) that they have to support. Thus the Tri-Laboratory community developed a new procurement model based on a common hardware environment over multiple years. However, to balance the risk associated with long term commitment to a single solution in the fast paced commodity space, we have chosen to limit the scope of the procurement to two government fiscal years with delivery across the fiscal boundaries (FY11-12).

By deploying a common hardware environment multiple times at all three laboratory sites over two government fiscal years, the time and cost associated with any one cluster is greatly reduced. In addition, it is anticipated that purchasing a huge set of common hardware components will lead to lower cost through volume price discounts. The Tri-Laboratory site splits for this procurement in Government Fiscal Years 2011-12 (4QCY10-3QCY12) have yet to be determined.

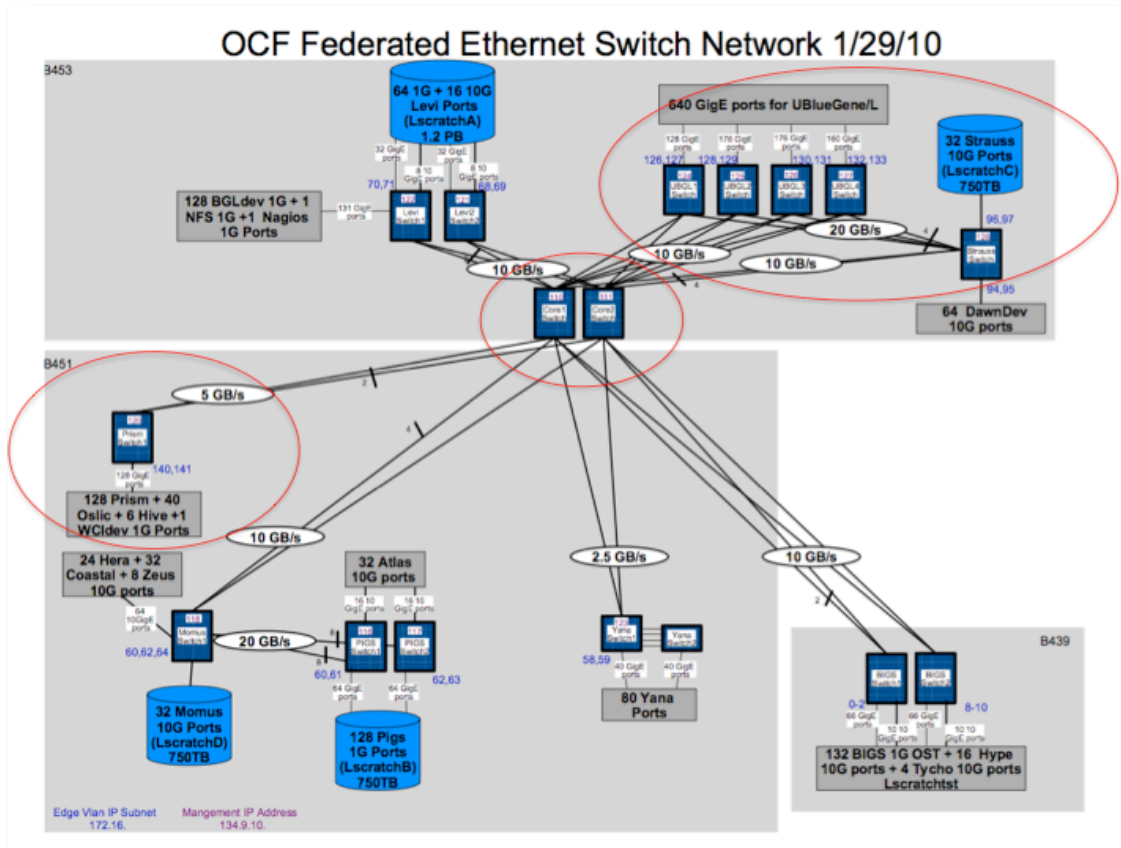
The ASC Program has also been successful in its use of the Tri-Laboratory Common Computing Environment (CCE) software. The ASC Program capacity systems strategy is to minimize the TCO of capacity systems through a common Tri-Laboratory hardware and software environment.

### **1.3 Tri-Laboratory Simulation Environments**

The ASC Program capacity systems will be deployed at the LLNL, LANL and SNL sites in the context of existing simulation environments. These simulation environments were developed as part of the ASC Program. Fundamentally, these simulation environments are built around a site-wide global file system (SWGFS) that is shared amongst all of the computing, visualization, networking and storage resources that comprise the simulation environment. LLNL and SNL have standardized their simulation environments on the Lustre parallel file system ([www.lustre.org](http://www.lustre.org)) and LANL has standardized its simulation environment on the Panasas parallel file system ([www.panasas.com](http://www.panasas.com)). Each Laboratory now has a significant investment (>\$100M) in their respective simulation environments.

#### **1.3.1 The LLNL Simulation Environment**

The LLNL Open Computing Facility (OCF) simulation environment is depicted in Figure 1. A similar simulation environment exists for classified computing (Secure Computing Facility or SCF). The elements of this simulation environment are on the floor at LLNL today. The simulation environment comprises five basic components: 1) 10-150 teraFLOP/s scale Linux clusters; 2) Federated 1 and 10 Gb/s Ethernet networking infrastructure; 3) scalable visualization resources; 4) HPSS based archival resources and 5) Lustre multi-cluster file system components. The Luster file system has three components: client, metadata servers (MDS) and object storage targets (OST). The Lustre client code runs on the compute and rendering nodes of the clusters. The Lustre MDS and OST components are comprised of commodity building blocks and RAID disk devices.



**Figure 1: The OCF simulation environment includes multiple 10-40 teraFLOP/s scale Linux clusters, visualization and archival resources. The unifying elements are a Federated 1 and 10 Gb/s Ethernet switching infrastructure (possibly moving to IB) and the Lustre file system.**

### 1.3.2 The LANL Simulation Environment

The bulk of the LANL simulation capability resides in the secure environment. The Linux portion of the LANL Secure Integrated Computing Network (ICN) is depicted below in Figure 2. Two more similar simulation environments exist for open computing: the Open ICN for unclassified open export controlled computing; and the Open Collaborative Network (OCN) for unclassified open collaborative computing. The Linux portion of the Secure ICN is in production today running important Weapons Program significant simulations day in and day out. The simulation environment comprises five basic components: 1) 50-1500 teraFLOP/s scale HPC platforms; 2) PaScalBB 900 GByte/sec Ethernet networking infrastructure; 3) scalable visualization resources; 4) HPSS based archival resources and 5) Panasas OBSD global parallel file system. The Panasas file system is a high performance parallel file system employing object based storage technology. The Panasas client code runs on the compute and rendering nodes of the clusters.

LANL currently makes heavy use of hardware-assisted GPU methods for its visualization needs, in particular its rendering and display needs. There is also a great deal of interest in

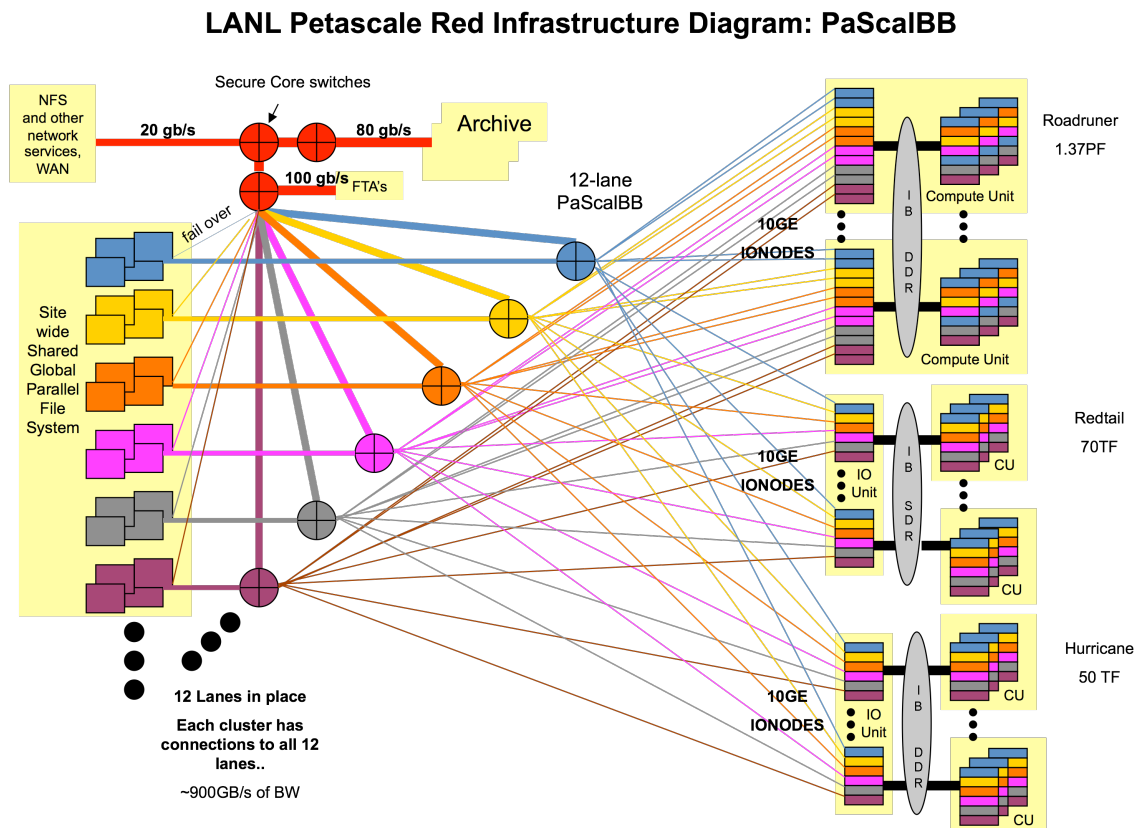
## TLCC2 Statement of Work

hybrid computing, specifically compute methods making use of GPU accelerators, augmenting the traditional CPUs.

For visualization, the rendering process of LANL's standard visualization tool uses GPUs for greatest efficiency. For its large-scale multi-panel CAVE and PowerWall stereo displays, LANL also uses high-end GPUs capable of swaplocking, genlocking and framelocking across several GPUs.

LANL's hybrid compute needs require GPUs with ECC and with fast double precision capability.

On the TLCC2 procurement, LANL is requesting the option of GPU enhanced nodes for a hybrid compute standalone cluster or sub-cluster. The GPUs chosen should satisfy the needs of GPGPU hybrid computing.



**Figure 2: The LANL Secure ICN simulation environment includes multiple 50-1500 teraFLOP/s scale HPC platforms, visualization and archival resources. The unifying elements are an innovative multipath resilient Parallel Scalable Ethernet Back Bone (PaScalBB) and the common global parallel Panasas file system (PanFS) based on the ANSI T10/1335-D Object Based Storage Device standard.**

### 1.3.3 The SNL Simulation Environment

The SNL Restricted Network (SRN) simulation environment is depicted in the Figure 3 below. A similar simulation environment exists for classified computing (Sandia Classified Network or SCN). The elements of this simulation environment are on the floor at the SNL California and New Mexico sites today. The simulation environment is comprised of five

basic components: 1) 3-325 teraFLOP/s scale Linux clusters; 2) 1 and 10 Gb/s Ethernet networking infrastructure; 3) scalable visualization resources; 4) HPSS based archival resources and 5) Lustre multi-cluster file system components. The Lustre file system has three components: client, metadata servers (MDS) and object storage targets (OST). The Lustre client code runs on the compute and rendering nodes of the clusters. The Lustre MDS and OST components are comprised of commodity building blocks and RAID disk devices. The figure below shows our Lustre environment connected to clusters, archive and visualization clusters.

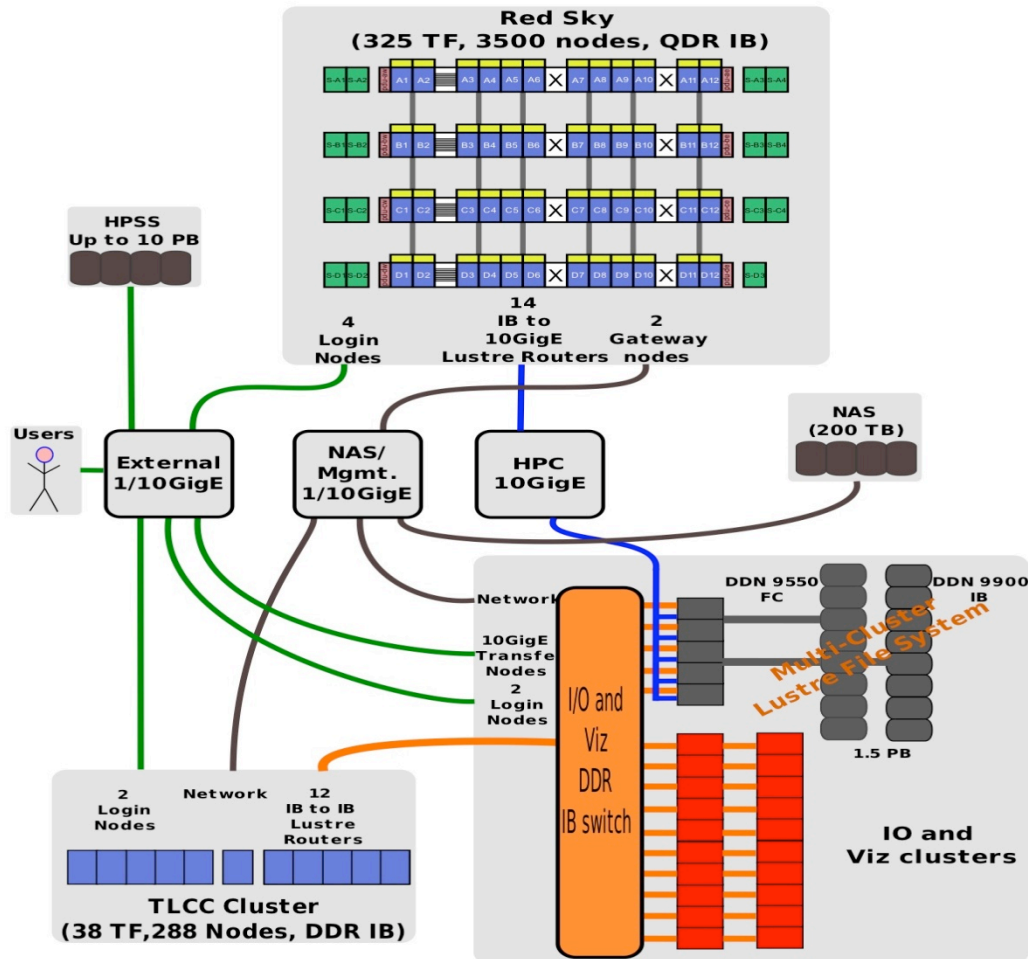


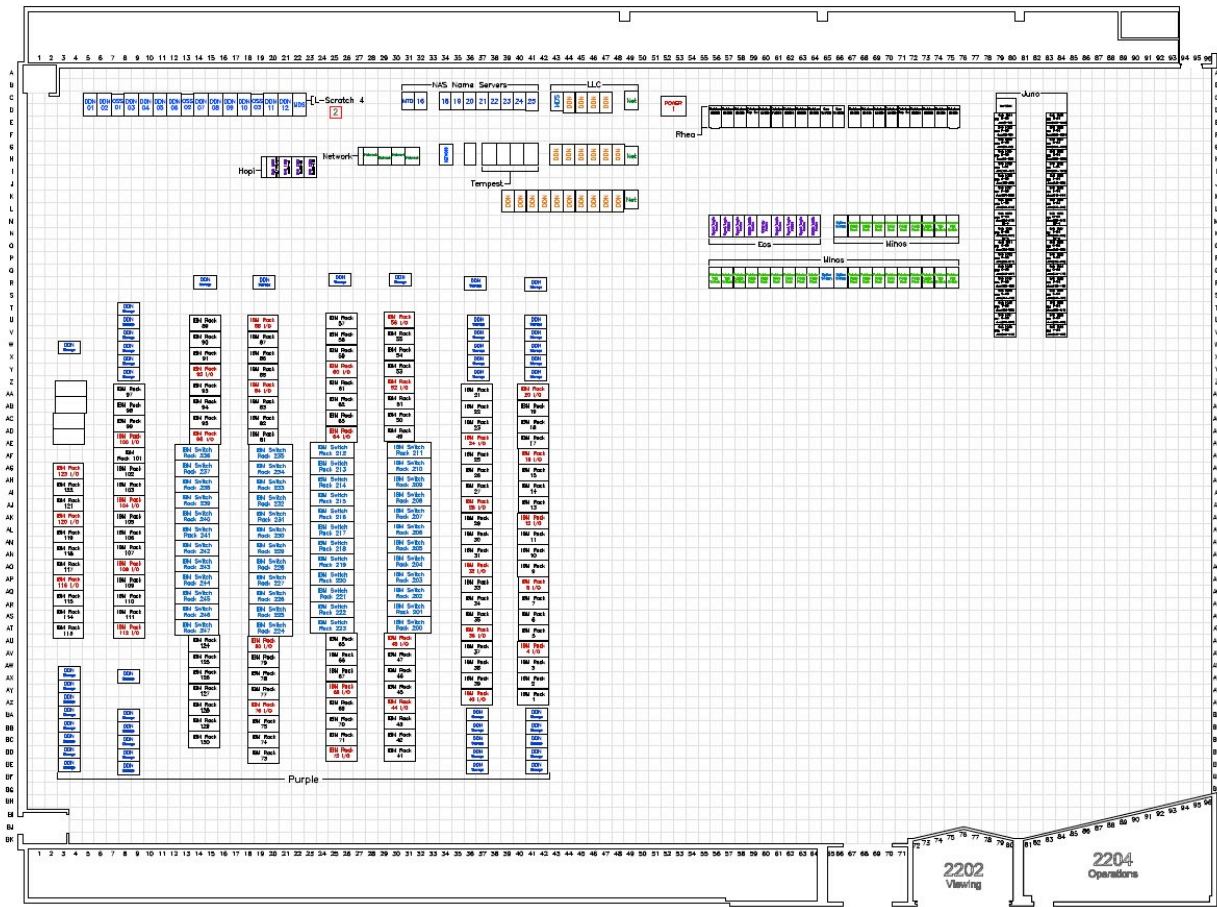
Figure 3: The SNL Simulation Environment is based on Lustre and Linux clusters

#### 1.4 Utilization of Existing Facilities

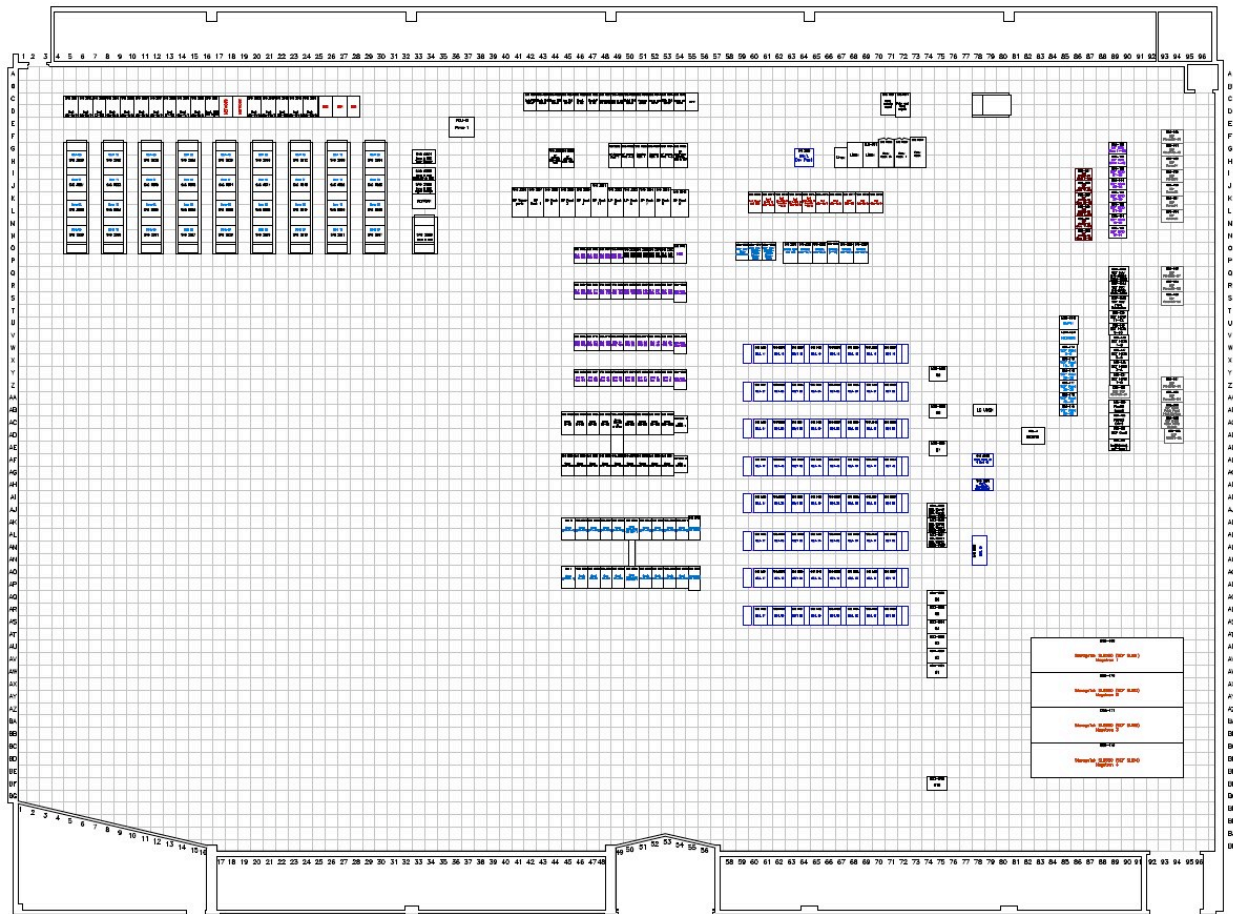
The SUs delivered under this subcontract will reside in the Terascale Simulation Facility (TSF) B453 and B451 at LLNL; the SCC facility Building 2327 and the LDCC facility building 1498 at LANL; and buildings/rooms 880/220, 880/230 and 725/105 at SNL Albuquerque and 912/097 at SNL Livermore. The following Sections provide high-level background information to Offerors on these facilities.



# TLCC2 Statement of Work



## TLCC2 Statement of Work



**Figure 4: Building 451 (top), Building 453 West (middle), and Building 453 East (bottom) computer floors are being reserved to site TLCC2 SU aggregations.**

Some TLCC2 SU aggregations will be migrated to classified operation with access to the SCF networking faculties. These SUs will be physically located inside a Limited Access Area in a Vault Type Room (VTR) in Building 453. LLNS will only provide access to the room to authorized personnel under Authorized Escort. All on-site personnel will be required to submit applications for access and be approved by standard LLNS procedures prior to entry into this facility. All on-site personnel will require being DOE P-cleared or P-clearable. RFP responses should indicate if the proposed on-site team has members that are non-U.S. citizens. Physical access to this facility by foreign nationals from sensitive countries will not be allowed. Dialup capability and Internet access to the system will be allowed up through acceptance. Authorized individuals may be allowed remote access for running diagnostics and problem resolution. Interaction of the on-site engineering staff with factory support personnel may be limited in some ways (e.g., dissemination of memory dumps from the system may be restricted).

An existing facility, B451, will also be used for siting the TLCC2 SU aggregations. See Figure 4. B451 facility has approximately 20,000ft<sup>2</sup> and 4 MW of power for computing systems and peripherals and associated cooling available for this purpose. The existing systems leave approximately 1.3 MW available for TLCC2 SUs. Facilities modifications to provide the necessary power and cooling for TLCC2 will need to be accomplished prior to SU delivery. It is therefore essential that the Offeror make available to LLNS detailed and **accurate** site requirements for the TLCC2 SUs at proposal submission time. The

## TLCC2 Statement of Work

Laboratory will be responsible for supplying the external elements of the power, cooling, and cable management systems.

These limitations emphasize the importance of local access to source code, particularly for operating system daemons.

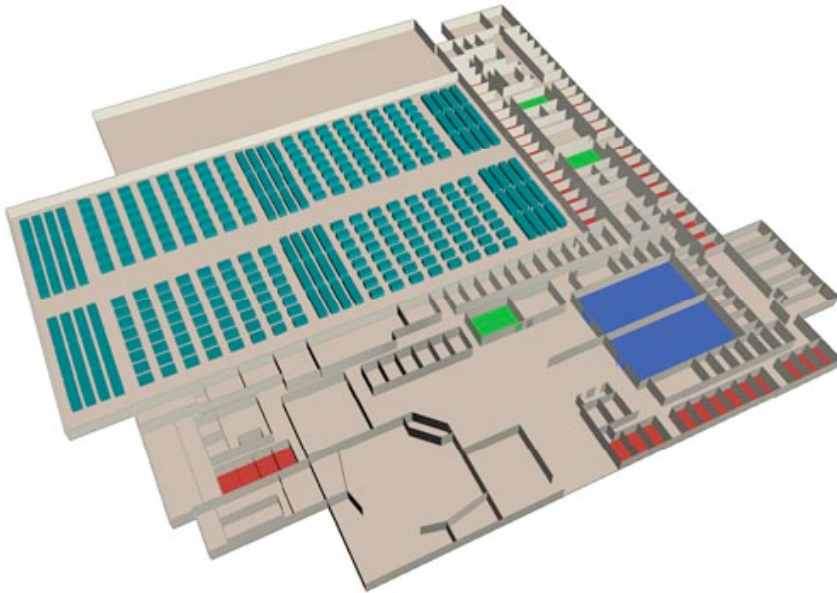
On-site space will be provided for personnel and equipment storage.

Personnel must practice safe work habits, especially in the areas of electrical and mechanical work.

### 1.4.2 The LANL Existing Facilities

Existing facilities, portions of Strategic Computing Complex (SCC) and LDCC computer floors, will be used for siting the TLCC2 SU aggregations. See Figure 5. The SCC will house the bulk of the SU aggregations in the LANL secure environment. The SCC has approximately 303,000ft<sup>2</sup> gross and 44,000ft<sup>2</sup> machine room space with 19.2 MW available for computing (out of approximately 30 MW building power) and 130,000 GPD expandable to 215,000 GPD water available for secure computational equipment. The LDCC has approximately 12,000ft<sup>2</sup> machine room space for open computational equipment. Facilities modifications to provide the necessary power and cooling for TLCC2 will need to be accomplished prior to SU delivery. It is therefore essential that the Offeror make available to the LANL detailed and **accurate** site requirements for the TLCC2 SU at proposal submission time. The LANL will be responsible for supplying the external elements of the power, cooling, and networking needed for system integration. Before those facility modifications, the available resources are:

	<b>LDCC (rm. 341)</b>	<b>SCC (main computer floor)</b>
Average airflow	600-800 CFM/floor tile	800-1500 CFM/floor tile
Electrical power distribution	208V/120V 3 phase PDUs	208V/120V 3 phase PDUs
Floor loading	250 lb/sqft	300 lb/sqft
Ceiling Height	10 feet	16 feet
Floor depth	24 inch	42 inch



**Figure 5: SCC Building 2327 diagram**

The SCC Building 2327 loading dock door where the semi backs up to is 10' high and 9' wide. Only LANL security guards can provide access for delivery. The semi driver will drop the trailer off and park the tractor outside the fence while the trailer is unloaded. It is ~70 feet to the computer room door from the edge of the loading dock. The computer room door is ~8' 10" high and 7' 8" wide. The computer room ceiling is ~16' high.

The LDCC Building 1498 has size restrictions on delivery trucks, which must fit under a ~13ft underpass and must back out without turning around. **IMPORTANT:** The doors leading to the computer floor limit the size of delivered equipment to at most 81" high and at most 65" wide (including delivery dollies). The freight elevator has 15,000 lbs payload capacity. The Offeror is responsible for safely moving computer racks from the truck to the computer room.

Some TLCC2 SU aggregations will be placed into classified operation with access to the Secure ICN facilities. In addition, TLCC2 SU aggregations will be physically located inside a Limited Access Area in a VTR. LANL will only provide access to the room to authorized personnel under Authorized Escort. All on-site personnel will be required to submit applications for access and be approved by standard LANL procedures prior to entry into this facility. All on-site personnel must be **U.S. citizens**. RFP responses should explicitly confirm that the proposed on-site team consists of U.S. citizens only, since physical access to this facility by foreign nationals will not be allowed. Dialup capability and Internet access to the system will be not be allowed. Interaction of the on-site engineering staff with factory support personnel may be limited in some ways (e.g., dissemination of memory dumps from the system may be restricted). These limitations emphasize the importance of local access to source code, particularly for operating system components.

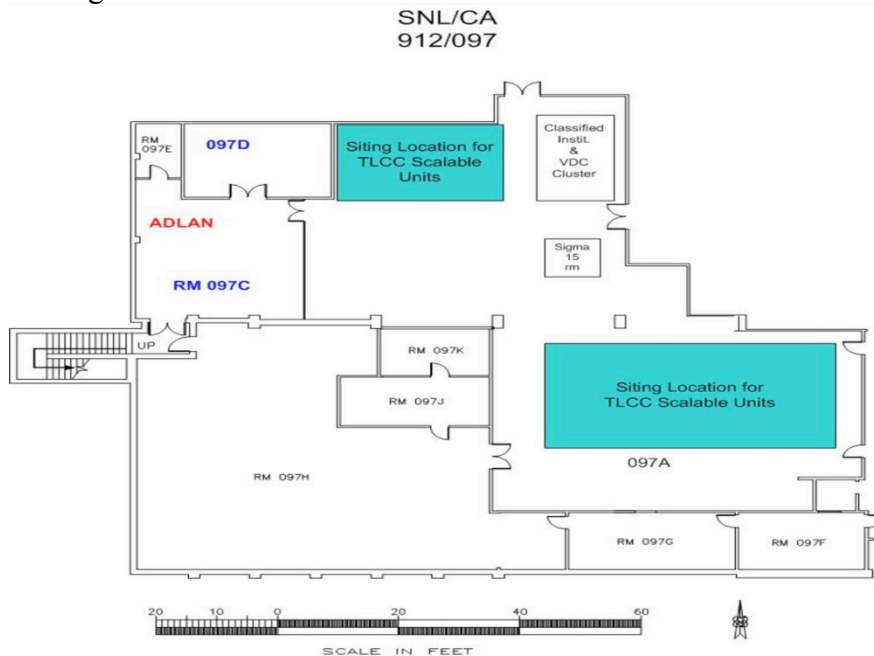
On-site space will be provided for personnel and equipment storage.

A safety plan will be required for on-site personnel. Personnel must practice safe work habits, especially in the areas of electrical and mechanical activities.

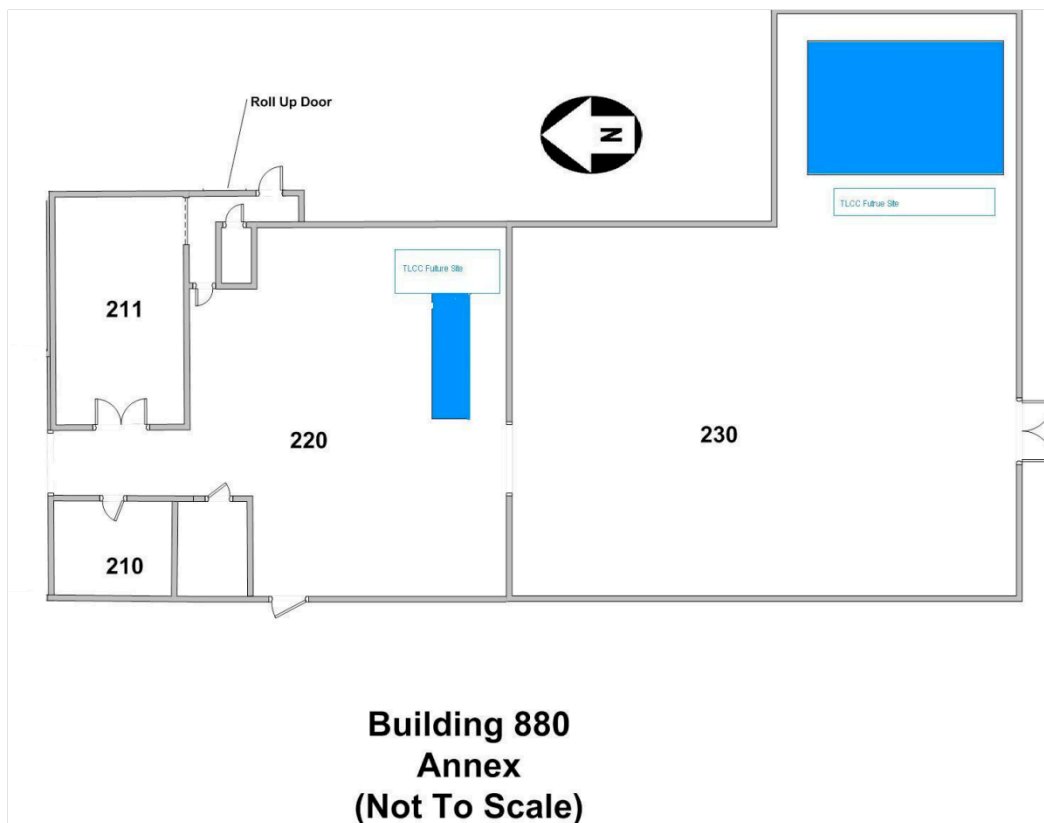
### 1.4.3 The Sandia Existing Facilities

The existing Sandia facilities, 912/097 at SNL California, and 880/230, 880/220 at Sandia New Mexico, can be used for the TLCC SU aggregations. See Figures 6 and 7. The SNL California building 912 facility has approximately 1,750 ft<sup>2</sup> (1,250+500 ft<sup>2</sup>) for the TLCC clusters. The 912/097 facility has 800 kW of power and 250 tons of cooling available for the TLCC clusters. The Sandia New Mexico 880/220 facility has approximately 300 ft<sup>2</sup>, 400 kW of power, and 130 tons of cooling. The Sandia New Mexico 880/230 room has ~ 500ft<sup>2</sup> available, 800 kW of power, and 500 tons of cooling. SU aggregation cannot span floors. Facilities modifications to provide the necessary power and cooling for TLCC will need to be accomplished prior to SU delivery (SNL may request a site visit to insure power and cooling requirements are per specification). It is therefore essential that the Offeror make requirements for the TLCC SU at proposal submission time. SNL will be responsible for supplying the external elements of the power and cooling.

The Figures of the SNL facilities are below.



**Figure 6: Sandia California Building 912 Computer Room 097**



**Figure 7: Sandia New Mexico Building 880 Computer Room 220/230**

In building 880 Computer Annex one semi-truck at time can back up to a multilevel dock height of 31" on the south side; on the north side dock height is 51", or unloading the truck with a fork lift would be fine; dock has a lift to take the computer racks down,(approx.3') to a dock door of 5' (wide) by 95" high to a slight 30 degree 12' ramp and then proceed 60' to an inside dock door (same dimensions) and then onto the computer room floor.

Some TLCC SU aggregations will be migrated to classified operation with access to the SCN networking faculties. In addition, TLCC SU aggregations will be physically located inside a Limited Access Area in a VTR. SNL will only provide access to the room to authorized personnel under Authorized Escort. All on-site personnel will be required to submit applications for access and be approved by standard Sandia procedures prior to entry into this facility. RFP responses should indicate if the proposed on-site team has members that are non-U.S. citizens. Physical access to this facility by foreign nationals from sensitive countries will not be allowed. Authorized individuals may be allowed remote access for running diagnostics and problem resolution. Interaction of the on-site engineering staff with factory support personnel may be limited in some ways (e.g., dissemination of memory dumps from the system may be restricted). These limitations emphasize the importance of local access to source code, particularly for operating system daemons. On-site space will be provided for equipment storage. A safety plan will be required for on-site personnel. All personnel entering the SNL computing facility must comply with all work control processes in place. Personnel must practice safe work habits, especially in the areas of electrical and mechanical activities.

## End of Section 1

## 2 TLCC2 Scalable Unit Strategy and Architecture

This section describes the overall TLCC2 Scalable Unit (SU) strategy and architecture.

### 2.1 TLCC2 Strategy

As described above, the Tri-Laboratory ASC Program community requires a large amount of capacity computing resources over the next Government fiscal year. In order to affordably and efficiently provide this Production quality computing capacity, the TLCC2 technical committee has embarked on an approach that extends existing practices within the Tri-Laboratory community while significantly improving TCO. During market survey discussions with industry, the Tri-Laboratory identified the approach of highly replicated “scalable units” that can be easily built, shipped, sited at the receiving laboratory and accepted quickly. In addition, there is a balance between the number of SU’s deployed and amount of work to maintain a large number of separate clusters at each site. Thus, the Tri-Laboratory has the need to aggregate SU’s into clusters built with 1, 2, 4, 8 and 16 SUs. The strategy is to purchase (under the subcontract resulting from this RFP) all of the components to build cluster SU aggregations and second stage InfiniBand Architecture (IBA) 4x QDR switches.

The delivered SUs in some set of aggregations called clusters will be integrated into existing classified simulation environments at the receiving laboratory. As such, the SU need to integrate into the receiving Laboratory’s existing multi-cluster file system. For LLNL and SNL, that multi-cluster parallel file system is Lustre from Oracle (<http://www.oracle.com>) and for LANL that multi-cluster parallel file system is PanFS from Panasas (<http://www.panasas.com/panfs.html>).

By replicating the SU many times during the subcontract, the Tri-Laboratory intends to reduce the cost to produce, deliver, install and accept each SU. In addition, this approach will produce a common Linux cluster hardware environment for the Tri-Laboratory user and system administration communities and thus reduce the cost of supporting Linux clusters and programmatically required applications on those clusters. However, the Tri-Laboratory prefers the SU design to be flexible enough to accommodate the following technology improvements:

1. Processor frequency improvements within the same cost and power envelopes
2. New processor socket and/or chipset improvements
3. New processor cores
4. Disks with higher capacity
5. New memory speed and capacity improvements
6. Interconnect bandwidth and latency improvements

Due to the extremely attractive cost/performance of x86 based Linux clusters and large number of x86 based clusters fielded at all Tri-Laboratory sites and the need for *at least* 2 GB of memory per processor core, this procurement focuses on solutions that are binary compatible with AMD x86-64 and Intel EM64T, and have InfiniBand 4x QDR interconnect or better. Dual socket nodes with multi-core processor implementations are widely available. The resulting B:F ratio, measuring the interconnect bandwidth of node (B) to the peak 64b floating point arithmetic performance of the node (F), has acceptable balance for our applications with this technology choice.

In order to minimize TLCC2 cluster support costs and the time to migrate a TLCC2 cluster into classified Production status, the Tri-Laboratory community will supply the Linux cluster software for building, burn-in and accepting the SUs. A Digital Versatile Disk (DVD) will be provided containing the TLCC2 CCE software stack, configured for this purpose. This variant of the CCE software stack consists of a RedHat Enterprise Linux distribution that has been enhanced to support vendor supplied hardware, cluster system management tools required to install, manage and monitor the SU, and a Tri-Lab workload test suite. It is the intent of the Tri-Laboratory community to use the InfiniBand software stack provided within Red Hat Enterprise Linux (RHEL) for use on Production computing clusters. Additional InfiniBand functionality may be added from the Open Fabrics Alliance or the greater open source community as needs arise. More specific details of CCE are provided in Section 2.3 below. The Tri-Lab workload test suite will be used as the SU burn-in, pre-ship test and then run again as a post-ship test after the SU is delivered and assembled at the receiving Laboratory. Once the SU is delivered, the Offeror and receiving Laboratory will be responsible for combining multiple SUs, as directed by the ASC Program user community, into clusters with the Offeror supplied spine switches and cables. Final acceptance of these clusters will be accomplished with a scaled up version of the pre/post-ship test.

The support model for TLCC2 is extremely simple. The Tri-Laboratory community will supply Level 1 and Level 2 support with the Offeror providing Level 3 support functions. On the software side, the Offeror is required to provide and support (level 3) device drivers and other low level specialized software for the provided hardware (including IBA). This support should be provided during normal working hours, Monday through Friday. On the hardware side, the Offeror is required to provide an on-site parts cache of Field Replaceable Units (FRUs) sufficient to cover a at least one week worth of failures without refresh and a Return Merchandise Authorization (RMA) mechanism for return of failed FRUs and refreshing the on-site parts cache. The Tri-Laboratory personnel at each site will provide Level 1 and Level 2 software and hardware support functions that includes responding to problem reports, root cause analysis, reading diagnostics and swapping FRUs. The Offeror shall provide training, on an as needed basis, for hardware FRU replacement over the lifetime of the support period of the subcontract. The Offeror will supply software maintenance for each Offeror supplied software component, starting with the first SU acceptance and ending three years after cluster acceptance.

## 2.2 TLCC2 Cluster Architecture

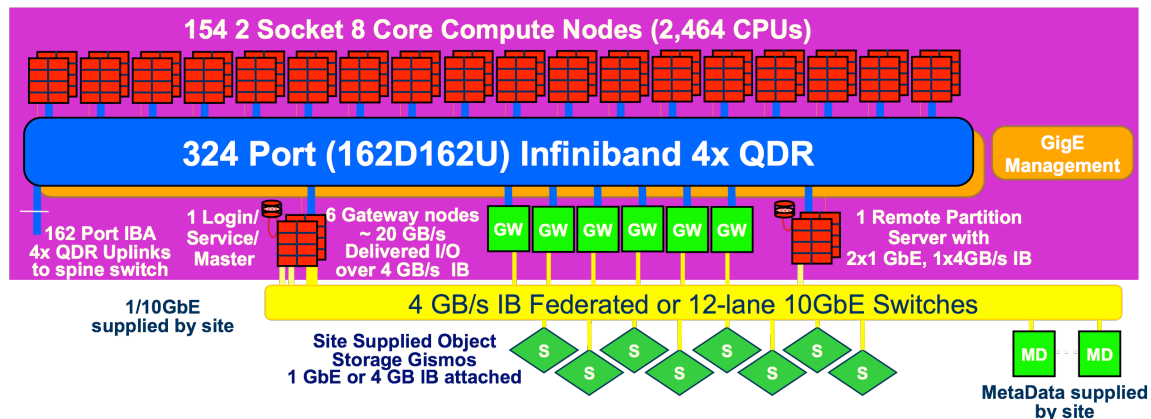
This section's description is illustrated by a specific, generic and vendor neutral SU point design. This point design is based on "generic" 2U white boxes with 324 port IBA 4x QDR switches. However, this choice for pedagogical purposes does not constitute a preference by the TLCC2 technical committee for this solution. The TLCC2 preference is for an optimized SU design that is more dense (including blades) than this example, yet still meets the facilities limitations for power, cooling and weight. That is, highly optimized and dense solutions (e.g., blades) that are architected in space wasteful ways are not seen as beneficial. For clusters larger than 4 SUs, the Offeror may choose to use small port count leaf switches (e.g., 36-port/single switch ASIC switches) to reduce SU costs, to ease node-to-switch integration, and to reduce the maximum number of hops within the IBA fabric.

The ASC Tri-Laboratory's scalable systems strategy for TLCC2 is almost the same as that previously implemented for TLCC07 (a prior procurement in GFY07). The basic idea is that

## TLCC2 Statement of Work

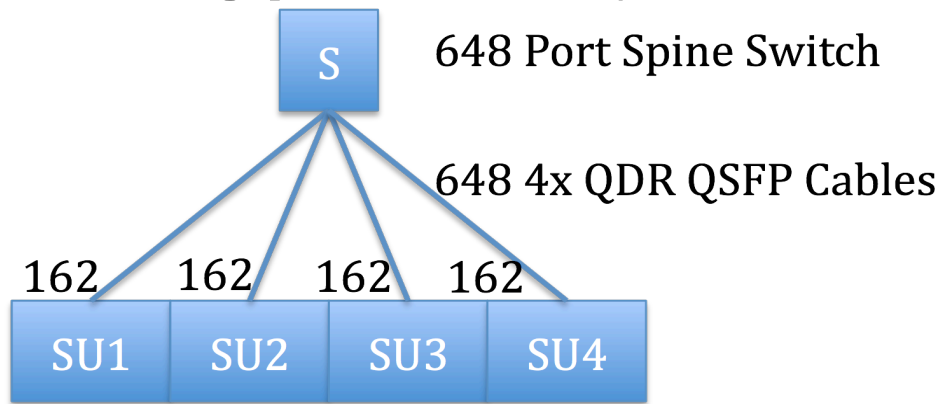
large clusters are usually built from two stage fat-tree interconnects with large radix switches. To architect a production quality cluster in a scalable fashion, the system resources like compute nodes, login nodes, management infrastructure and IO infrastructure are divided into smaller groupings associated with each first stage switch. The exact number of each component depends on the size of the switch and the capacity and bandwidth requirements for various components. Thus, a large cluster can be scaled up in these replicated SUs. When contemplating purchasing a large number of clusters of various sizes over a one-year period, this approach allows structuring the acquisition and integration activity into a very large number of replicated SUs. This provides the Offeror numerous opportunities to optimize and parallelize SU component purchases, building, testing, shipping, installation and accepting activities.

One difference from TLCC07 is that, in TLCC2, LANL desires an option for a special SU where compute nodes are enhanced with or replaced by GPU-enabled nodes capable of handling a variety of hybrid computing workloads. LANL anticipates a need for one or two of these special SU's for delivery during this subcontract. Neither SNL, nor LLNL desire these GPU enhanced SUs.

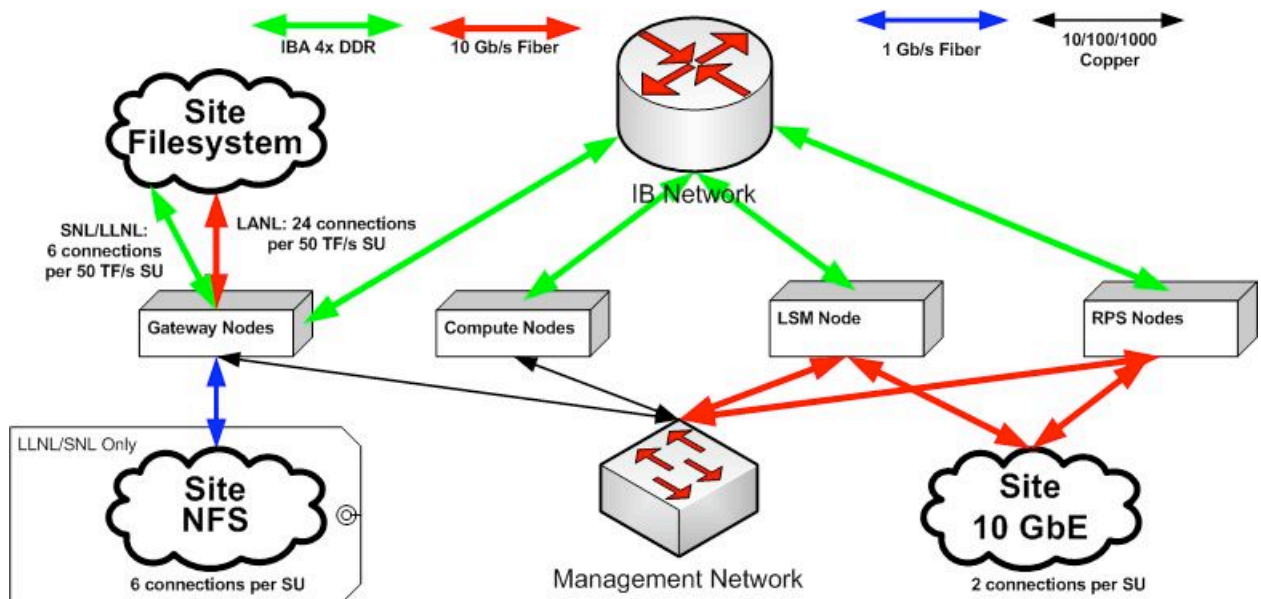


**Figure 8: Example TLCC2 Scalable Unit architecture.** Large clusters can be built up by aggregating various numbers of these SU. TLCC2 cluster architecture includes clustered I/O model, no local node disks, dedicated login/service/master nodes, dedicated gateway nodes and compute nodes all connected to site supplied networking and attached RAID disk resources for Lustre or PanFS.

$$648 \times 40 \text{ gbps} \times 2 = 6,480 \text{ GB/s}$$



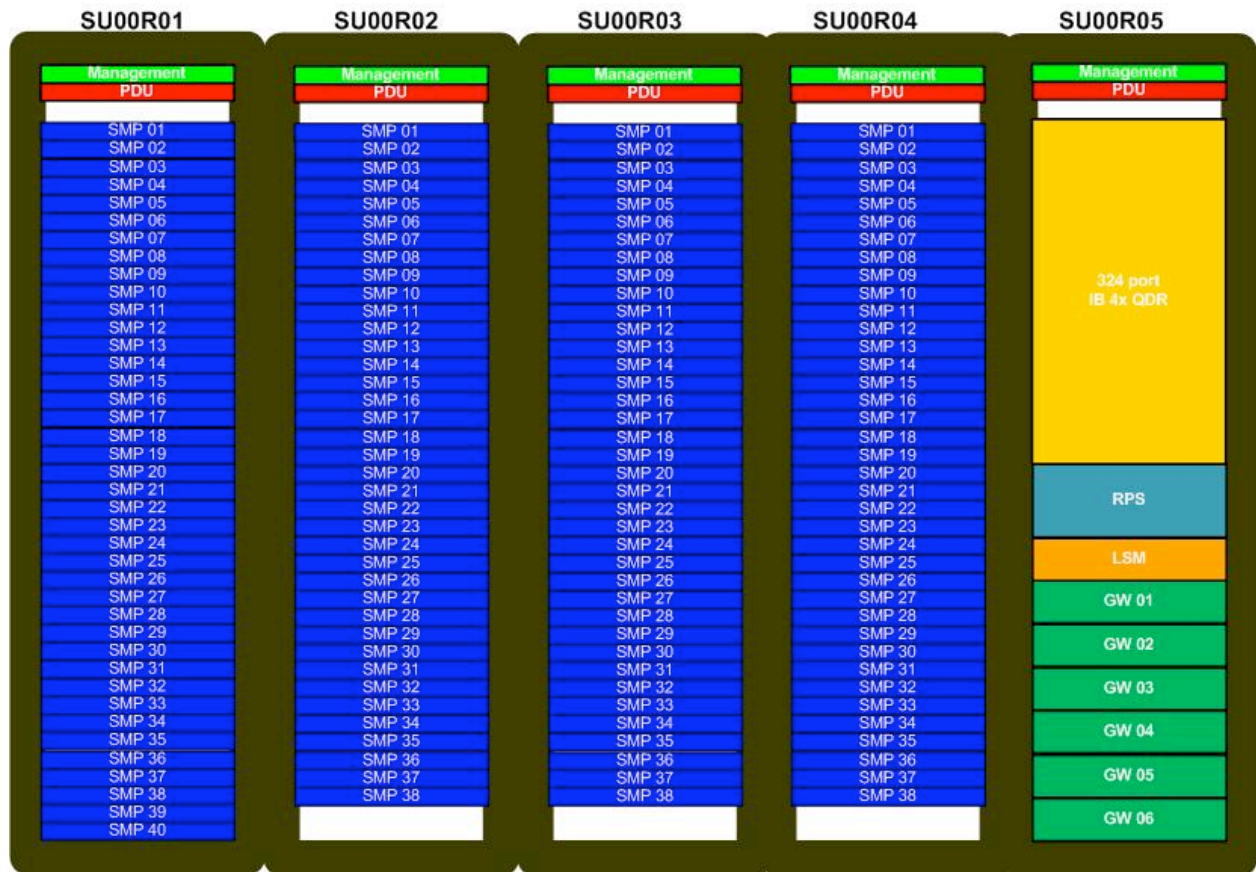
**Figure 9: Infiniband 4x QDR interconnect for TLCC2 SU is based on the port counts of the 324 and 648-port IBA 4x QDR switches. Clusters deployed by Tri-Laboratories will be aggregations of multiple SU configured with a multi-stage, full bandwidth, non-blocking, fat-tree federated IBA switch. This 4xSU example is based on a single 648 4x QDR switch.**



**Figure 10: TLCC2 Network Layout**

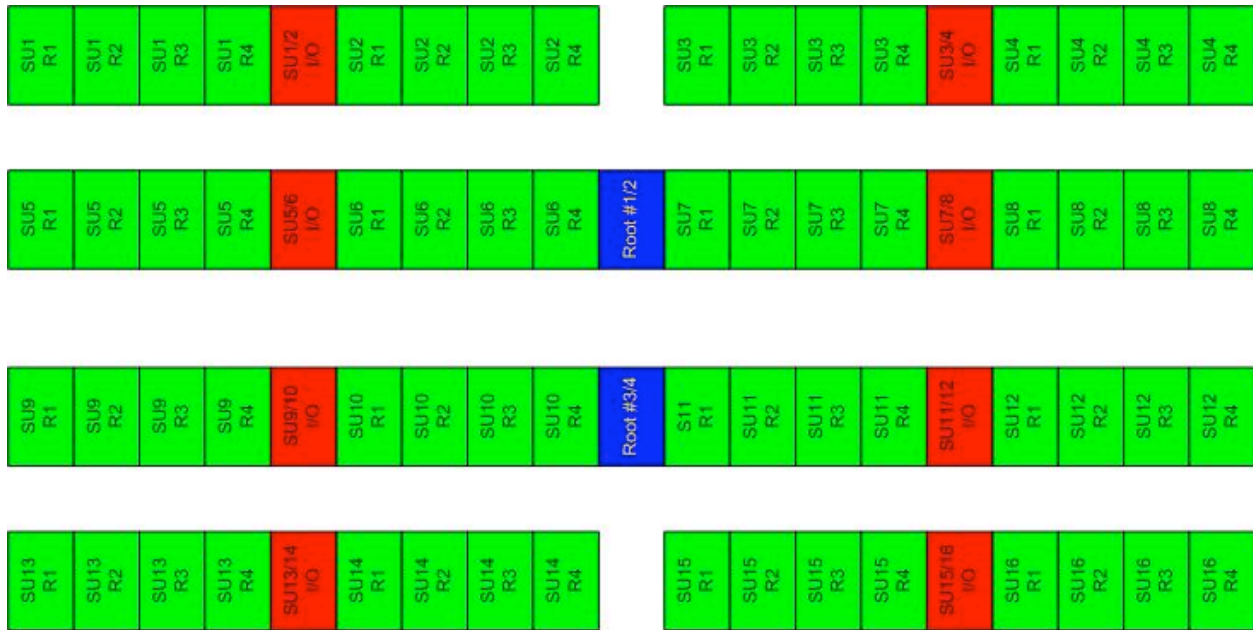
The 50 TF/s, 5 TB memory SU example in Figure 8 is based on 162 2U/2-Socket nodes with a single port of 4x QDR IB over PCIe2 x8. The SU has 20 GB/s global I/O bandwidth using either six 4x QDR IB or 24 x 10GigE through the six gateway nodes. This SU requires **5x42U** compute racks for compute nodes and terminal server 1x42U IO rack for a combination of compute nodes, gateways, login/service/master and remote partition server nodes, terminal server and the management Ethernet switch.

The 5x42U rack SU example in Figure 11 is about 10' long and 3' wide. Air flows from the front of the row to the back. Within the SU, the distances between all the nodes and the IBA switch are much less than 10 meters. The compute nodes each require about 500 watts and hence each compute rack is about 20 KW and weighs about 800 lbs or about 133 lbs/ft<sup>2</sup>.



**Figure 11: TLCC2 SU five rack layout based on 154 1U/2 Socket nodes plus larger nodes for GW (6), LSM (1) and RPS (1)**

After the SUs are delivered, installed and accepted, the receiving Laboratory and the Offeror will combine multiple SUs together to form a TLCC2 Cluster. The floor plan layout in Figure 12 shows a hypothetical installation of a 16 SU cluster. Each set of four SUs is in a separate row with the 648-port IBA 4x QDR spine switch in Red rack the middle of the row. The two SUs on the left of the spine switch rack have the rack layout shown in Figure 11. The two SUs on the right of the spine switch have the rack layout reversed (mirror image). With a slight re-arrangement of the GW nodes (i.e., moving them to the compute racks), the 2 SUs can share rack 5 and thus the pair of SUs require just 9 racks. The blue racks depict top level 648-port switches to make one cluster.



**Figure 12: Example TLCC2 16 SU Cluster Layout.** Each row is based on 4 SUs with 4 1/2 racks per SU including the IBA 4x QDR Spine Switch. The root of the IB network consists of four 648 port central switches (in blue).

With this layout, cable distances between the SU IBA switch and the IBA spine switches are less than 10m. Each row of four SUs has four login/service nodes each with a connection to the IB network, two 1 Gb/s Ethernet and one 40 Gb/s InfiniBand (or one 10 Gb/s Ethernet card with two ports of 10 GbE) connection to the Laboratory infrastructure. In addition, each row has twenty-four gateway nodes each with either one 4x QDR IB InfiniBand, or four 10Gb Ethernet connections to the Laboratory infrastructure.

Note in Figure 12 that the cold isles are wider than the hot isles and the airflow through the racks emanates from floor tile grates in front of the racks in the cold isle. Air exhausts from the racks into the hot isle and is removed from the room via grates in the ceilings above the hot isles. There are no air handlers on the floor at LLNL or at the LANL SCC facility, but will be present at SNL and the LANL LDCC facility. In addition, power is provided to the racks by cables running under floor from wall panels; otherwise, overhead cabling is **required**. To minimize facilities modifications costs and breaker utilization, TLCC2 racks should be designed to use a minimum number of circuits (one) and maximize the utilization of that circuit (up to the maximum of 80%).

## 2.3 TLCC2 Software Environment

To execute the ASC Program capacity systems strategy, the TLCC2 SUs must be integrated into each site's infrastructure and transitioned to Production service as quickly after acceptance as possible. The software required to do so is the Tri-Laboratory Common Computing Environment (CCE), built from Red Hat Enterprise Linux distribution and additional 3<sup>rd</sup> party and open source software.

The CCE is a set of software components common among the Tri-Laboratory community. The CCE activity started in parallel with TLCC07 procurement and continues today. Each

receiving Laboratory will install the production CCE software on each SU after acceptance, configured for local system integration. This software will be targeted to the Offeror's hardware environment in collaboration with the Offeror after subcontract award and prior to the SU manufacture, so that it may support pre-ship testing, acceptance testing, and eventually production deployment of TLCC2 clusters at each site.

All CCE software and components are self-supported. Tri-Laboratory Open Source developers work closely with system administrators and users to resolve problems on production systems. For any given software package, there is a designated package owner who handles release, test as well as any support issues that arise in production. Depending on the nature of the package, owners may be the primary developer and fix bugs themselves, or they may be the liaison to an external support resource.

External support relationships are primarily developer-to-developer. In the case of Red Hat, the Tri-Laboratory community has access to a full-time Red Hat engineer who works directly with TLCC2 systems and support people and acts as the liaison to Red Hat for everything in the Red Hat Linux distribution.

For the purposes of the TLCC2 test and acceptance activity, the Tri-Laboratory community will create a **single** variant of the CCE software stack, including the Tri-Laboratory Operating System Stack (TOSS). This single CCE variant combines elements of the CCE software and is focused on providing the Offeror with the software that will be needed to build, debug and validate each SU, along with the suite of applications that will be required for acceptance. This software will be tailored to the Offeror's hardware environment in collaboration with the Offeror after subcontract award and prior to the SU manufacture.

Section 2.3.1 below outlines the TLCC2 test plan, followed by a description of CCE (including TOSS) and any significant differences between the sites. It is important for the Offeror to understand the eventual software environments that will be deployed by the Tri-Laboratories once the SU are accepted. Over time, those software environments will evolve while maintaining commonality of key hardware and software components between sites (see section 1.2).

### **2.3.1 TLCC2 Synthetic Workload Test Plan**

For each SU or SU aggregation, pre-ship and post-ship acceptance tests will be conducted in three major phases: functionality test (~1 day), performance test (~1 day), and stability test phase (5 days) during which Synthetic Workload (SWL) test suite will be run repeatedly. Details of the testing protocol will be specified in the TLCC2 SWL Test Plan to be negotiated with the successful offeror.

In the past TLCC07 procurement, the functionality test phase included hardware configuration, system administration, software configuration, and other functionality tests. Additionally, this test phase included MPI validation test suite, MPI-I/O Romio tests, Moab/SLURM functionality tests, OpenMP microbenchmark, TOSS QA test suite, and testing of typical system administration functions.

The performance testing phase included Presta MPI performance tests, NEWS05 asynchronous communication stability and performance test, STREAM and STRIDE

memory performance tests, TTCP and NetPerf network tests, MATMUL matrix-matrix multiply tests, as well as MPI-Bench collectives. Additionally, either Lustre or Panasas serial and parallel tests (as appropriate for the installation site) may be run during acceptance tests. Sites using Lustre would also conduct LNet self-test.

The stability test phase content includes HPL (Linpack), hydrodynamics codes (sPPM, Miranda, Raptor), radiation codes (UMT, IRS, SWEEP3D), molecular dynamics (SPaSM, LAMMPS), plasma (YF3D/Yorick), quantum chromodynamics (QCD), solvers (AMG, Trilinos-Eptra) and various benchmarks (MPI, NPB, IOR, HPCCG). Automated submission of SWL workload through Gazebo tool allows for continuous 24hr/day operation during stability tests, with subsequent collection and analysis of results stored on the RPS node.

### **2.3.2 Description of Common Computing Environment (CCE)**

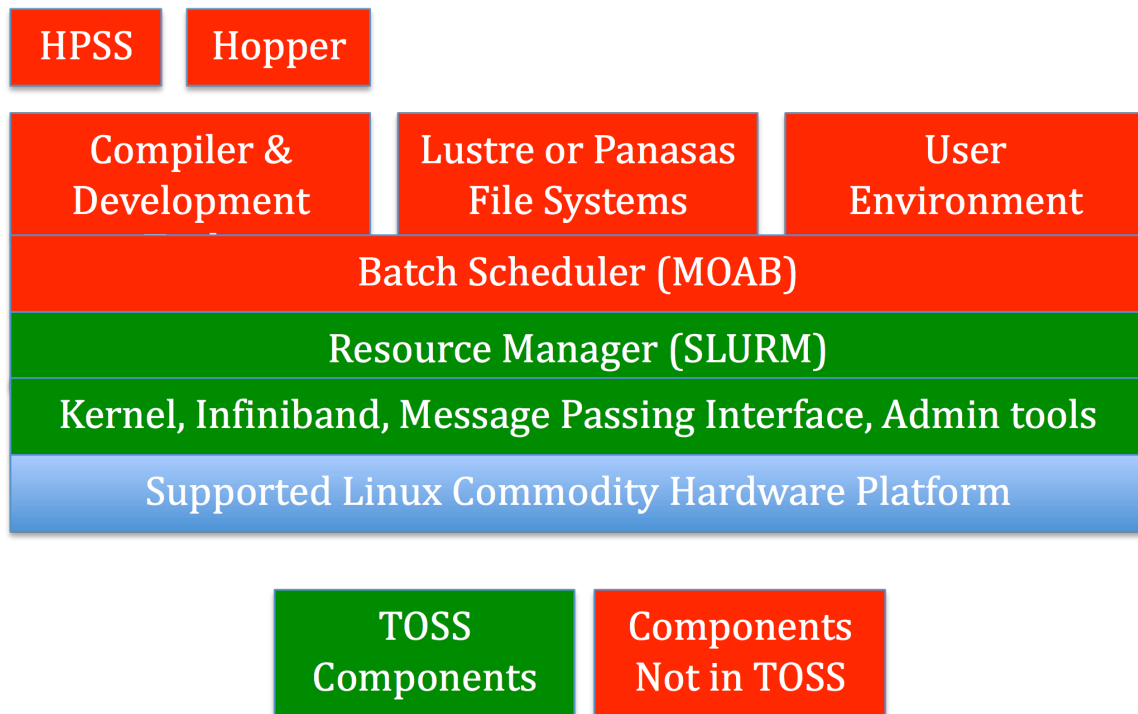
As stated earlier, the Tri-Laboratory community produces and supports a software environment for HPC Linux clusters called the CCE, based on the TOSS, built from RHEL distribution and additional software components.

All CCE software and components are self-supported. Tri-Laboratory Open Source developers work closely with system administrators and users to resolve problems on production systems. For any given software package, there is a designated package owner who handles release, test as well as any support issues that arise in production. Depending on the nature of the package, owners may be the primary developer and fix bugs themselves, or they may be the liaison to external support resources, such as 3<sup>rd</sup> party software vendors.

External support relationships are primarily developer-to-developer. In the case of Red Hat, the Tri-Laboratory community has access to a full-time Red Hat engineer who works directly with TLCC2 systems and support people and acts as the liaison to Red Hat for everything in the Red Hat Linux distribution.

As of March 2010, CCE is based on TOSS 1.3 and scheduled for deployment on existing TLCC07 platforms Tri-Lab wide. Some key components of this 2010 software stack are RHEL 5.4, TOSS 1.3 patches and tools (including SLURM 2.1), Gazebo 1.1 test framework, Open MPI 1.3.3, mvapich 0.9.9, Lustre 1.8.2, Perceus 1.5.0, OneSIS 2.0.1, NFSroot 2.16, cfengine 2.2.1, Open Fabrics InfiniBand Software, environment modules 3.1.6, plus 3<sup>rd</sup> party software such as TotalView debugger, production compilers (Intel 11.0.081, Pathscale 3.1, PGI 8.0.1), Moab scheduler, and Panasas client. These components will be updated by the time of TLCC2 subcontract award.

At the core of TOSS 2.0 (scheduled for deployment in 2011) will be the Red Hat Enterprise Linux 6.x (RHEL6) distribution, including a number of additional cluster-aware components. Some components of the RHEL distribution are modified to meet the demands of high-performance computing installations, operations and support. Additional separately licensed 3<sup>rd</sup> party components such as PanFS client (where deployed), compilers, TotalView debugger, and Moab scheduler are not considered a part of TOSS, yet they complete the CCE production environment.



**Figure 13: TOSS**

### 2.3.2.1 Tri-Laboratory Operating System Stack (TOSS)

The TOSS distribution contains a set of RPM (Red Hat Package Manager) files, RPM lists for each type of node (compute, management, gateway, and login), and a methodology for installing and administering clusters. It is produced internally and therefore supports a short list of hardware and software. This approach permits each site to support a large number of similar clusters with a single TOSS release, supported by a small staff, and to be agile in planning its content and direction.

Fundamental components of TOSS are:

- A complete RHEL distribution augmented as required to support targeted HPC (e.g., TLCC2) hardware and cluster computing in general.
- A RHEL kernel that is optimized and hardened to support large scale cluster computing, including EDAC3 support for all TLCC2 platforms.
- The InfiniBand software stack including MVAPICH and OpenMPI libraries, and Subnet Manager (SM) scalable to full size TLCC2 systems.
- The SLURM resource manager with support for both MVAPICH and OpenMPI over InfiniBand, full NUMA awareness, single job scalability to 78,848 processors, and a compatibility library to support TORQUE job submission command syntax.
- Fully integrated Lustre and Panasas parallel file system clients and Lustre server software. Panasas client is separately licensed.
- Scalable cluster administration tools to facilitate installation, configuration (including BIOS setup/upgrade), and remote lights-out management.
- An extensible cluster monitoring solution with support for both in-band and out-of-band (e.g., IPMI) methods.
- A PAM authentication framework for OTP and Kerberos and an access control interface to SLURM.

## TLCC2 Statement of Work

- A test framework for hardware and operating system validation and regression testing, extensible to include Tri-Lab tests.
- GNU C, C++ and Fortran90 compilers integrated with MVAPICH and OpenMPI.

TOSS forms the foundation, upon which additional software components of CCE are layered.

Additional comments on TOSS components, some of which need to be provided by the Offeror:

Distribution – It is expected that by the time of the TLCC2 contract award, the CCE software stack will be derived from RHEL version 6 or later.

Kernel – TOSS replaces the RHEL kernel with an enhanced kernel. This kernel includes additions in the areas of device support for InfiniBand, VFS modifications for Lustre, ECC and FLASH memory device support for Intel motherboard chipsets, crash dump support, miscellaneous bug fixes, and optimized configurations for TLCC2 hardware.

InfiniBand (IB) Stack – The RHEL Open Source OpenFabrics Enterprise Distribution (OFED) IBA software stack will be provided as part of TOSS. In addition, the OpenSM IB InfiniBand subnet manager as well as other Open Fabrics and open source InfiniBand extensions are provided. The Offeror is expected to use OpenSM as the fabric subnet manager throughout the acceptance test stage. The Tri-Labs will work with Offeror to incorporate IB extensions into TOSS if required for IB functionality. This additional IB functionality shall be accessible using open source tools.

Device drivers – Any device drivers required to support the Offeror's hardware, which are not available in the RHEL standard distribution or need enhancements, should be provided by the Offeror for incorporation into TOSS. This additional or modified software must be provided in the form of buildable source RPMs with licensing terms, which allow for the free redistribution of that source (BSD or GPL preferred).

Diskless cluster installation and configuration – The TOSS stack includes tools for diskless cluster installation and configuration, such as NFSroot, Perceus, and OneSIS.

Remote Management - The Offeror should provide node firmware required to implement the IPMI 2.0 protocol. FreeIPMI will be used to validate basic IPMI 1.5 and IPMI 2.0 compliance. PowerMan (<http://sourceforge.net/projects/powerman/>) will be provided for remote power management. ConMan (<http://home.gna.org/conman/>) will be provided for remote console management.

Cluster Monitoring – Ganglia and/or SNMP may be used to gather data of interest from the nodes (in-band). This data may include node resource utilization such as CPU, memory, I/O, runaway processes, etc. Node environmental data (e.g., temperature, fan speeds, voltages) may be collected by utilizing the LMSensors kernel module or with FreeIPMI's ipmi-sensors running out-of-band.

Resource Manager –Simple Linux Utility for Resource Management (SLURM, see <https://computing.llnl.gov/linux/slurm/>).

Compilers – TOSS comes with GNU, Intel, Pathscale and PGI compiler suites, but their functionality may be limited by licensing requirements. A temporary license for at least one major 3rd party compiler may be provided during pre-ship tests.

Test Harness – Gazebo (LANL) uses Moab scheduler to submit test workload, and then collects test outputs for analysis.

Synthetic WorkLoad (SWL) – Set of applications representative of Tri-Laboratory workload used with Gazebo test harness to stress test the SU and clusters of SU aggregations. This SWL will only contain unclassified codes that are not export controlled.

BIOS management tools - Any BIOS management tools required to support the Offeror's hardware should be provided and maintained by the Offeror, for incorporation into TOSS.

**Firmware** — Firmware images for standard motherboards, including FLASH/CMOS support software, is included in TOSS. Firmware and support software for power control/serial console hardware is also included. *Any additional firmware required by the proposed hardware should be supplied and maintained by the Offeror.*

**System Initialization** — [NFSroot](#), Perceus, or OneSIS are used to provision images to diskless nodes.

**YACI** — Yet Another Cluster Installer is Livermore's system installation tool based on various cluster installers such as VA system imager and LUI.YACI can fully install the 1,152-node MCR cluster in about 15 minutes. It is image-based and can use either an NFS pull or multicast mechanism to install many nodes in parallel.

**Genders** — Genders (<http://sourceforge.net/projects/genders>) is a static system configuration database and rdist Distfile preprocessor. Each node has a list of "attributes" that in combination describe the configuration of the node. The genders system enables identical scripts to perform different functions depending on their context. An rdist Distfile preprocessor expands attribute macros into node lists allowing very concise Distfiles to represent many large clusters.

**Cfengine** — Configuration engine is a widely used framework to configure and fully proscribe every aspect of a platform (<http://www.cfengine.org>). Cfengine is used to configure and ensure software state convergence on HPC platforms at LANL.

**ConMan** — The ConMan console manager (<http://home.gna.org/conman/>) manages serial consoles connected either to hardwired serial ports or remote terminal servers (telnet based), performs logging of console output, and manages interactive sessions, permitting console sharing, console stealing, console broadcast, and interfaces for transmitting a serial break or resetting a node via PowerMan. Conman will support IPMI through FreeIPMI's Ipmiconsole utility and libipmiconsole library.

**PowerMan** — The PowerMan power manager

(<http://sourceforge.net/projects/powerman/>) manages system power controllers and is capable of sequenced power on/off for groups of nodes and initiating reset (both plug off/on and hardware reset if available). Powerman currently supports IPMI 1.5 or 2.0 through FreeIPMI's ipmipower and other hardware. It can be extended to support new hardware.

**Host Monitoring System** — TLCC2 may monitor in-band (while Linux is running) by polling via NET-SNMP (<http://net-snmp.sourceforge.net/>). Among the information polled is motherboard sensor information and information about failing hardware devices such as memory and disks. In addition, PowerMan can extract out-of-band monitoring information such as case temperature from some remote power control devices that have this capability. LLNL's SNMP based host monitoring system stores current state in a MySQL database and long-term state in an RRD (round robin database). Collection software polls cluster nodes in parallel using SNMP bulk queries and a sliding window algorithm to reduce polling latency. Status is presented via web using Apache and PHP. LANL's host monitoring uses external solution based on Zenoss (<http://www.zenoss.org>).

**FreeIPMI** - FreeIPMI (<http://www.gnu.org/software/freeipmi>) provides in-band and out-of-band IPMI software based on the IPMI v1.5/2.0 specification. FreeIPMI supports various IPMI subsystems including sensor monitoring, system event log (SEL) monitoring, power management, chassis management, watchdog, serial-over-LAN (SOL), and a number of OEM extensions.

### 2.3.2.2 Simple Linux Utility for Resource Management

SLURM is an Open Source, fault-tolerant and highly scalable cluster management and job scheduling system for clusters containing thousands of nodes. SLURM is the production resource manager on all Tri-Lab TLCC07 clusters, and it has been ported to other systems as well (<https://computing.llnl.gov/linux/slurm/slurm.html>). SLURM is a part of TOSS.

The primary functions of SLURM are:

- Monitoring the state of nodes in the cluster.
- Logically organizing the nodes into partitions with flexible parameters.
- Accepting job requests.
- Allocating both node and interconnect resources to jobs.
- Monitoring the state of running jobs, including resource utilization rates.

While SLURM can support a simple queuing algorithm, Moab Cluster Suite will manage the order of job initiations through its sophisticated algorithms described in Section 2.3.3 of this document.

SLURM utilizes a plug-in authentication mechanism that currently supports authd and the LLNL-developed munge protocol. The design also includes a scalable, general-purpose communications infrastructure. APIs support all functions for ease of integration with external schedulers. SLURM is written in the C language, with a GNU autoconf configuration engine. SLURM's modular design allows for ease of portability.

### **2.3.3 Moab Scheduler**

Moab Cluster Suite is a professional cluster management solution that integrates scheduling, managing, monitoring and reporting of cluster workloads. Moab is a separately licensed 3<sup>rd</sup> party software package not included in TOSS. Moab Cluster Suite simplifies and unifies management across one or multiple hardware, operating system, storage, network license and resource manager environments to increase the ROI of cluster investments. Its task-oriented graphical management and flexible policy capabilities provide an intelligent management layer that guarantees service levels, speedy job processing and easily accommodates additional resources. For more information see <http://www.adaptivecomputing.com/products/moab-hpc.php>

### **2.3.4 Lustre Cluster Wide File System (Sandia, LLNL)**

Sandia and LLNL utilize the Lustre Cluster Wide File System on clusters built up from the TLCC2 SU. Currently, Lustre is in production status with TLCC07 SU's. Lustre hardware is not a part of TLCC2 procurement.

### **2.3.5 Panasas PanFS Multi-Cluster Environment Wide Global Parallel OBSD Based File System (LANL)**

LANL uses the centralized and globally shared Panasas File System (PanFS). See [www.panasas.com](http://www.panasas.com) for more information on PanFS. Panasas hardware is not a part of TLCC2 procurement.

Currently, LANL has three PanFS global parallel file systems in production, one in each computing environment (secure, open and collaboration networks). At LANL, the TLCC2 systems would be added into the PaScalBB network shown in Figure 14.

### LANL Petascale Red Infrastructure Diagram: PaScaBB

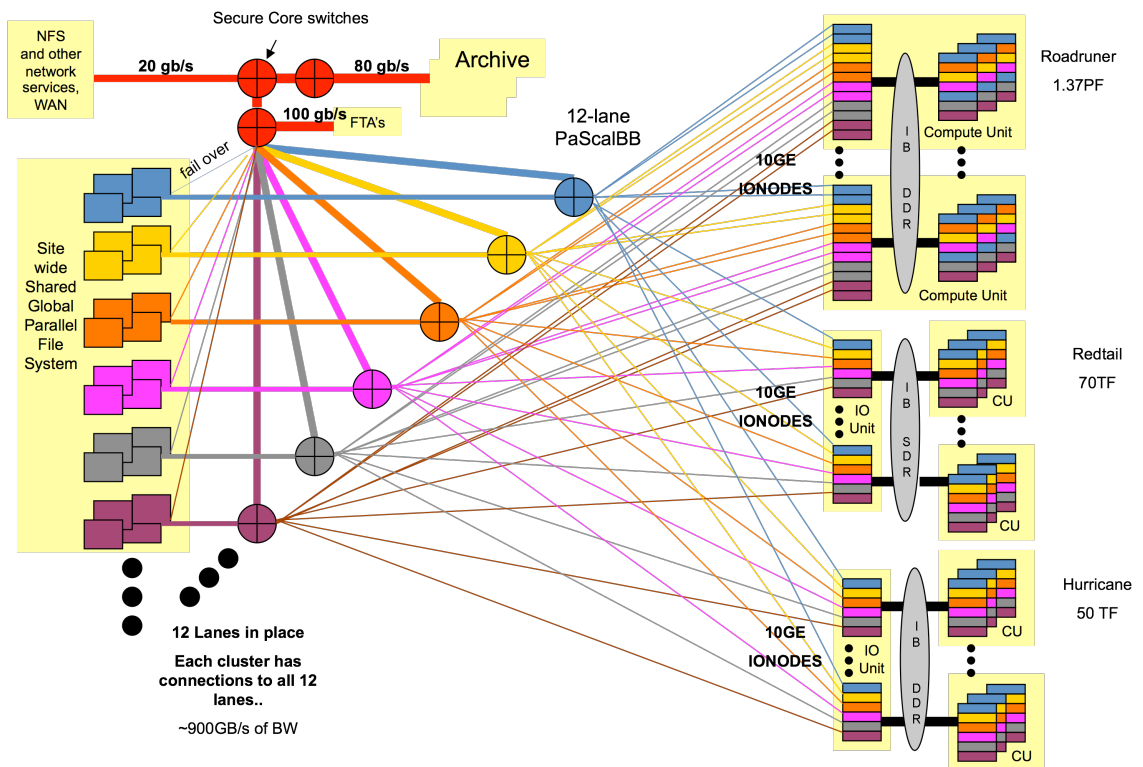


Figure 14: Secure Linux Environment at LANL. TLCC2 SUs would be incorporated as additional clusters.

#### 2.3.6 Production Compiler Suites

CCE includes Fortran, C and C++ compiler suites from Intel, PGI and PathScale. These 3<sup>rd</sup> party products are licensed separately and are not a part of TOSS. Tri-Labs *may* provide a temporary license for a production compiler suite, so that pre-ship tests can be rebuilt as necessary.

#### 2.3.7 TotalView Debugger

CCE includes the TotalView debugger. This 3<sup>rd</sup> party product is licensed separately and is not a part of TOSS. Tri-Labs do *not* expect to license TotalView before TLCC2 systems are deployed on each site.

End of Section 2

### 3 TLCC2 Technical Requirements

The end product of the TLCC2 procurement is a set of highly integrated, well-balanced capacity compute SUs with at least 50 TF/s, but not more than 250 TF/s as depicted in the SU example in Figure 8. Each SU will have compute, IBA interconnect, gateway, remote partition, and login/service/master resources. These SUs must be combinable in aggregations of at least 1, 2, 4, 8, or 16 SUs to form fully functional “capacity” clusters. The successful Offeror will be responsible for building, passing pre-ship testing with Tri-Laboratory software, delivering, installing, and passing post-ship testing of individual SUs. The successful Offeror, with the receiving Laboratory, will integrate SUs into integrated, fully functional clusters and pass cluster acceptance testing. The successful Offeror will work with the Tri-Laboratory Linux cluster community to integrate necessary device drivers and IBA software into the TOSS Linux distributions (see section 2.3). As directed by LLNS, the Offeror will provide aggregations beyond 4 SUs with sufficient additional IBA switches and cables to allow the Tri-Laboratory and the Offeror to construct clusters with full bandwidth, non-blocking IBA interconnects. These combined SUs shall be capable of supporting a complex workload consisting of small (4-256) medium (257–2,048), large (2,049-16,384) and occasionally full capability (78,848) MPI task count parallel jobs for Tri-Laboratory classified ASC Program and SSP simulations. TLCC2 SUs will reliably run production scientific simulations of a wide number of physical phenomena of importance to all SSP Campaigns and Directed Stockpile Work (DSW). The fully functional SUs and clusters comprised of aggregations of multiple SUs must be useful in the sense of being able to deliver a large fraction of peak performance to a diverse scientific and engineering workload. In particular, the SUs and clusters comprised of aggregations of multiple SUs must be capable of running a single user application with one MPI task per core over all compute nodes in the cluster. The SUs and clusters comprised of aggregations of multiple SUs must also be useful in the sense that the code development and production environments are robust and facilitate the dynamic workload requirements. They must also be easy to install, manage and operate in order to lower the Tri-Laboratory TCO.

To satisfy these demanding requirements, we anticipate needing a large set of tightly coupled SUs that integrate with Lustre or PanFS global file systems through high-speed external 4x QDR InfiniBand networking, or external 12-lane 10 Gb/s Ethernet infrastructure at LANL. Our requirement is to have these SU built from commodity AMD x86-64 or Intel EM64T (or binary equivalent) nodes containing at least two (2) microprocessor sockets. These SUs shall have IBA 4x QDR (or faster) compatible interconnect consisting of IBA switches, cables, and adapters. In addition, these SU shall have 1 Gb/s Ethernet and a second 4x QDR InfiniBand connection for external networking, or multi-lane PaScalBB 10 Gb/s Ethernet external networking at LANL.

This subcontract will be structured with deliveries commencing in 3QCY11 and ending in 1QCY12. During this period of time, there may be advancements in COTS technology utilized in any proposed SU configuration. As such, the Offeror shall provide these technology enhancements to the Tri-Laboratory community in future quarterly deliveries of SU, and may offer to upgrade previously delivered SU as separately priced options. The Offeror shall state which technology enhancements are expected to be delivered and the circumstances required to trigger those enhancements.

Mandatory Requirements (designated MR) in the Draft Statement of Work (SOW) are performance features that are essential to Tri-Laboratory requirements. An Offeror must

satisfactorily propose all Mandatory Requirements in order to have its proposal considered responsive.

Mandatory Option Requirement (designated MOR) in the Draft SOW reflects a particular Scalable Unit (SU) configuration required by LANL. LANL needs the ability to acquire this SU configuration as an option. An Offeror must satisfactorily propose all MOR in order to have its proposal considered eligible for award of a subcontract for LANL SUs.

Target Requirements (designated TR-1, TR-2, or TR-3), identified throughout the Draft SOW, are features, components, performance characteristics, or other properties that are important to the Tri-Laboratory. However, omission of a response for a Target Requirement will not render a proposal non-responsive. Target Requirements add value to a proposal. Target Requirements are prioritized by dash number. TR-1 is most desirable to the Tri-Laboratory, while TR-2 is more desirable than TR-3. Target Requirement responses will be considered as part of the proposal evaluation process.

A listing of technical MRs, MORs, and TRs is included in the Draft SOW Table of Contents.

In addition to MRs, MORs, and TRs identified in this Draft SOW, the Offeror may choose to propose any additional features (i.e., Offeror proposed features) consistent with the objectives of the TLCC2 procurement and the Offeror's project plan, which the Offeror believes will be of value to the Tri-Laboratory. MRs, MOR, TRs, and additional features proposed by the successful Offeror, and of value to the Tri-Laboratory, will be stated as firm requirements in a final negotiated Statement of Work and incorporated in the resulting TLCC2 Subcontract.

### **3.1 High-Level Hardware Summary (TR-1)**

Offeror will provide a high-level overview of the proposed SU design (section 3.1.1) and its evolution (section 3.1.3) over the 3QCY11 through 1QCY12 timeframe. The intent of this section is to have in one place a technical summary of the Offeror's proposed SU deliveries. It is vital that the Offeror make absolutely clear in the response to these subsections, what will be delivered and when.

#### **3.1.1 SU High-Level Architecture**

Offeror's response to this section will contain a detailed description of the proposed TLCC2 SU and the proposed evolution of this SU technology over time. The features and functionality of all major components of the SU shall be discussed in detail. The Offeror will provide an architectural diagram of the TLCC2 SU, similar to Figure 8, labeling all component elements and providing bandwidth and latency characteristics (speeds and feeds) of and between elements. The Offeror will provide an architectural block diagram for each TLCC2 node type bid, labeling all component elements and providing bandwidth and latency characteristics (speeds and feeds) of and between elements. The node architectural diagrams will specifically show and label the chipset used and denote independent PCIe buses and slots and label these with bus widths and speeds. The Offeror will provide an architectural block diagram of the proposed IBA interconnect for the SU and for combining SUs in at least 1, 2, 4, 8 and 16 multiples similar to Figure 9. Offeror will provide a rack layout diagram for the proposed SU similar to Figure 11 and floor layout for at least four clusters consisting of aggregations of four SU each, similar to Figure 12. If Offeror proposes to deliver different SUs packaging configurations with differing

rack layouts in order to meet site specific power, cooling requirements (see section 3.2.12 and subsections), then a rack layout diagram for each proposed SU packaging configuration will be provided. Any alternative cooling strategies with non-trivial facilities impacts should be described, including liquid cooling preferred for the LDCC facility at LANL.

### 3.1.2 SU Requirements Summary Matrix

The following matrix identifies the highest priority technical requirements (TR-1) and will be completed in its entirety. Entries shall be labeled N/A if the requirement is not offered. In addition, the system requirements summary matrix will be completed for any alternate proposed systems submitted.

Index	Requirement Description	Qty	Offeror Response
1.	Compute node product designation		
2.	Compute node form factor		
3.	Compute node processor type, speed, and cache sizes		
4.	Compute node SPECfp2006 and SPECfp2006_rate		
5.	Compute node memory bus type and speed		
6.	Compute node chip set designation		
7.	Compute node number of expansion busses and types		
8.	Compute node number and type of expansion bus slots for each bus		
9.	Type and size of compute node memory		
10.	Compute node blade-chassis type and configuration, if applicable		
11.	GPU-node card product designation		
12.	GPU-node number of GPU cards per node		
13.	GPU-node type and size of node memory		
14.	LSM node product designation		
15.	LSM node processor type, speed and cache sizes		
16.	LSM node memory bus type and speed		
17.	LSM node chip set designation		

TLCC2 Statement of Work

Index	Requirement Description	Qty	Offeror Response
18.	LSM node number of expansion busses and types		
19.	LSM node number and type of expansion bus slots for each bus		
20.	Type and size of LSM node memory		
21.	Type and size of LSM node local SATA disk		
22.	Gateway node product designation		
23.	Gateway node processor type, speed and cache sizes		
24.	Gateway node memory bus type and speed		
25.	Gateway node chip set designation		
26.	Gateway node number of expansion busses and types		
27.	Gateway node number and type of expansion bus slots for each bus		
28.	Number and type of each PCIe expansion card(s) installed in each Gateway		
29.	Type and size of gateway node memory		
30.	RPS node product designation		
31.	RPS node processor type, speed and cache sizes		
32.	RPS node memory bus type and speed		
33.	RPS node chip set designation		
34.	RPS node number of expansion busses and types		
35.	RPS node number and type of expansion bus slots for each bus		
36.	Type and size of RPS node memory		
37.	Type and size of RPS disks and RAID config. Indicate RAID packaging solution (e.g., internal to node, external expansion chassis)		

Index	Requirement Description	Qty	Offeror Response
38.	Number and type of each PCIe expansion card(s) installed in each Gateway		
39.	RAID controller designation and interface types and numbers		

### 3.1.3 SU Evolution Overview (TR-1)

The Tri-Laboratory requires that the SU that are aggregated into a specific cluster at any site be as identical as possible. However, the Tri-Laboratory also requires that when processor, interconnects, memory and disk technology elements advance during the lifetime of the subcontract resulting from this procurement, that these enhancements will be integrated into future SU deliveries without perturbing SU cost or reliability significantly. Offerors will describe the anticipated technology advances and the circumstances required to trigger their integration into future SU deliveries. Offerors need not propose to upgrade SU hardware after delivery. Offeror will offer at least the following technology enhancements:

1. Processor frequency improvements within the same cost and power envelopes
2. New processor socket and/or chipset improvements
3. New processor cores
4. Disks with higher capacity
5. Higher speed and capacity memory improvements
6. Interconnect bandwidth and latency improvements

Offeror will provide at least the following information for these technology improvements. Overall SU Impact should be rated as low, medium or high and then major components that are impacted will be listed. Offeror will use “low” impact designation to indicate that no other major components are impacted by the change. Offeror will use “medium” impact designation to indicate that other major components of the SU require update, but not a new design and the SU architecture does not change substantially. Offeror will use “high” impact designation to indicate that other major components of the SU require redesign and/or the SU architecture does change substantially.

Item	Item Upgrade	Delivery Qtr	Attribute	Overall SU Impact
Processor	Speed Bump	4QCY11	X.X GHz clock	Low
		1QCY12	X.X GHz clock	Low
Processor	New socket		Socket Type, clock	Medium, new motherboard, new memory type/speed
Processor	Next Generation Processor	4QCY11	Processor name, socket, GHz clock, power	High, new motherboard, node design, new memory type/speed, new node design

Item	Item Upgrade	Delivery Qtr	Attribute	Overall SU Impact
Processor	OTHER			
Memory	DDR3 or FBD	4QCY11	More Bandwidth	Medium, new motherboard
Memory	OTHER			
IBA	Other			
Local Disk	HDD capacity bump	4QCY11	4 TB	Low

For proposed technology improvements that have medium or high impact to SU architecture design, Offeror will provide high-level SU architectural diagrams defined in section 3.1.1 for each.

## 3.2 SU Hardware Requirements

For each of the following SU hardware requirements, provide information for the first SU installation only. Changes to the proposed SU hardware to meet these requirements over the subcontract timeframe are covered in section 3.1.3.

### 3.2.1 TLCC2 Scalable Unit (MR)

Each SU the Offeror provides shall be based on at least two socket AMD64 or Intel EM64T (or binary compatible) nodes. All nodes in SUs to be aggregated into a specific cluster at any site shall be of the same processor and chipset revision. There shall be four node types on the IBA switch: compute nodes; gateway nodes; login/service/master node; and remote partition server. All nodes shall be attached to the management Ethernet network and also the internal IBA network. Depending on the deployment site, the login/service/master, RPS and gateway nodes shall also attach to the site-supplied 40 Gb/s Fiber 4x QDR InfiniBand infrastructure, or to multi-lane PaScalBB 10 Gb/s Ethernet infrastructure (at LANL). The login/service/master nodes shall also attach to a 1 Gb/s Fiber Ethernet infrastructure. The remote partition server shall be attached to the management Ethernet with 1 Gb/s copper Ethernet and also have 1 Gb/s fiber Ethernet infrastructure.

### 3.2.2 TLCC2 SMP two Socket Configuration (TR-1)

The SU compute and login/service/master nodes the Offeror provides will be based on at least two socket AMD64 or Intel EM64T (or binary compatible) nodes. If required to meet the node performance requirements below, the gateway and remote partition server nodes may also be four sockets.

### 3.2.3 SU Peak (TR-1)

It is desirable to size SUs reasonably, so that each SU the Offeror provides will have a peak 64b floating point processing power of at least 50.0 teraFLOP/s ( $50 \times 10^{12}$ ). In the following sections "F" in the B:F ratios denotes the Offeror's proposed node or SU 64b floating-point peak rate.

### 3.2.4 Number of SU (MR)

N, the number of SU required over the lifetime of the subcontract, is 2,150teraFLOP/s divided by the SU peak teraFLOP/s (delivered in 3QCY2011) rounded to the nearest integer.

Example: a 50 teraFLOP/s peak SU yields  $N = \text{round}(2,150/50) = 43$ . A 100 teraFLOP/s peak SU yields  $N = \text{round}(2,150/100) = 22$ .

Offerors shall provide at least 2,150 teraFLOP/s in overall aggregate performance over at least 22 SU.

### **3.2.5 TLCC2 Node Requirements**

The following requirements apply to all node types except where superseded in subsequent sections.

#### **3.2.5.1 Processor and Cache (TR-1)**

The SU nodes will be configured with at least two AMD64 or Intel EM64T (or binary compatible) multi-core processor dies with at least 1 MB of L2 or L3 on-chip cache. Each socket will utilize less than 105 Watts per processor socket when running a copy of Linpack on every core and be one speed grade slower than the fastest available in that wattage category available when the SU is built.

#### **3.2.5.2 Node Floating Point Performance (TR-1)**

The SU node will have performance on the SPECfp2006 of at least 240 . Offeror should detail both the SPECfp2006 and the SPECfp2006\_rate performance numbers.

#### **3.2.5.3 Chipset and Memory Interface (TR-2)**

The SU nodes will be configured with DDR3-1600 (800 MHz) or faster SDRAM. Assuming 4 memory channels per processor socket, the aggregate peak memory bandwidth is 51.2 GB/s/socket. Assuming a two socket node, the aggregate peak memory bandwidth becomes 102.4 GB/s/node, and  $162 \times 102.4 = 16.6$  TB/s/SU. The SU memory configuration will be architected and configured to deliver low memory latency and high bandwidth to ASC applications. The node will be configured with at least two PCIe2 (or faster) busses with one slot per bus. It is highly preferred that the chipset for all nodes in the cluster be the same.

#### **3.2.5.4 Node Delivered Memory Performance (TR-1)**

The SU nodes will be configured to deliver at least 30 GB/s per processor socket of memory bandwidth when running one copy of the Streams benchmark per core. Offeror will report with proposal the Streams performance running one copy of the streams benchmark per core, for each bid TLCC2 node type. See [https://asc.llnl.gov/computing\\_resources/purple/archive/benchmarks/](https://asc.llnl.gov/computing_resources/purple/archive/benchmarks/) or <http://www.cs.virginia.edu/stream/> for the streams benchmark.

#### **3.2.5.5 Node I/O Configuration (TR-2)**

The SU nodes chip set will be configured with two independent PCIe2 8x (or faster) buses and one PCIe2 slot per bus. The SU node PCIe infrastructure will be fully compatible with both the proposed IBA HCA and the proposed 10 GbE NIC.

#### **3.2.5.6 Node Memory (TR-1)**

The SU nodes will be configured with at least 2.0 GB memory per processor core. The

memory should be DDR3-1600 (800 MHz) or faster SDRAM with registered DIMMs, ECC and chipkill. As an example, this would result in 64.0 GB (4 GB DIMMs) for 4-socket 8-core processor solutions and 32.0 GB for 2-socket 8-core processor solutions.

### **3.2.5.7 Node Power (TR-2)**

The SU nodes components together will utilize less than 200 Watts per socket plus 150 Watts when running a copy of Linpack on every core and one speed grade slower than the fastest available in that wattage category available when the SU is built. Related to this specification, internal power supplies will meet at minimum the 80 PLUS Bronze level criteria, as specified at: [http://www.80plus.org/manu/psu/psu\\_join.aspx](http://www.80plus.org/manu/psu/psu_join.aspx), which is as follows: minimum efficiencies of 82%, 85%, and 82% at rated loads of 20%, 50% and 100% respectively. Power supplies will maintain a true power factor of 0.9 or greater at 50% rate load and higher. The test procedure for such measurements may be found at <http://www.efficientpowersupplies.org>. Unrated power supplies (of any output) can be sent to <http://80plus.org> and publicly posted to the website for a testing and reporting fee. Fee inquiries addressed on a case-by-case basis with 80 PLUS program staff.

### **3.2.5.8 Linux Access to Memory Error Detection (TR-1)**

The SU nodes will include a hardware mechanism to detect correctable and uncorrectable memory errors from Linux. This hardware mechanism will be capable of sending a non-maskable interrupt (NMI) or machine check exception when an uncorrectable error occurs, so that software may take immediate action. When a correctable or uncorrectable memory error occurs, this hardware mechanism will provide sufficient information so that software may identify the affected failed or failing DIMM FRU (i.e. the exact DIMM location on the motherboard). Node memory controller will expose defective memory modules that the chipkill or other functionality compensates for to the operating system down to the DIMM FRU level. Documentation that is sufficiently well written and complete so that Tri-Laboratory personnel can actually program this hardware facility (at the Linux level) will be delivered with the SU.

### **3.2.5.9 No Local Hard Disk (TR-1)**

The SU compute and gateway nodes will be configured **without** a traditional rotating disk-based hard drive.

### **3.2.5.10 Node Form Factor (TR-2)**

The Offeror will provide SU compute nodes with equivalent form factor of not more than 2U.

Denser (including blades) solutions, meeting the power and cooling requirements in section 3.2.12 and facilities requirements in section 0, are desired.

### **3.2.5.11 Integrated Management Ethernet (TR-1)**

The TLCC2 SU nodes will have at least one integrated copper 1 Gb/s management Ethernet.

### **3.2.5.12 Node BIOS (MR)**

Offeror shall provide all nodes of an SU with a fully functional BIOS that shall take a node from power on or reset state to the start of the Linux kernel as loaded from a

network connection or local disk (if a local hard disk is installed). The requirements below apply to all nodes types except where called out as specific to a particular node type.

#### **3.2.5.12.1 Node BIOS Type Options (MR)**

Offeror shall provide complete BIOS documentation for the delivered SU nodes. This documentation shall include description of all parameters, default factory settings and how to modify them, use and abuse of the tools to manage BIOS. If Offeror provides a proprietary, Offeror supported BIOS, a complete set of Linux-based command line tools shall be provided to manage BIOS.

#### **3.2.5.12.2 Remote Network Boot (MR)**

All provided node BIOS shall be capable of booting the node from a remote OS image across the Ethernet Management network. Booting over Ethernet shall utilize PXE, BOOTP, DHCP, or other public protocols.

#### **3.2.5.12.3 Remote IBA Network Boot (TR-1)**

All provided node BIOS may be capable of booting the node from a remote OS image across the IBA interconnect network. Booting over IBA may utilize an IBA standard protocol like PXE, SRP, iSER, IPoIB, or similar.

#### **3.2.5.12.4 Node Initialization (TR-1)**

The node BIOS initialization process will complete without human intervention (e.g., no F1 keystroke to continue) or fail with an error message written to the console. The time required for a node's BIOS to take the node from a power on state (or reset state) to the start of the loading of the Linux kernel boot image shall be less than thirty (30) seconds. Note that this includes the POST phase configured with minimal POST hardware checks. Shorter times are clearly desirable and achievable.

#### **3.2.5.12.5 Error/Interrupt Handling (MR)**

The delivered BIOS shall not block reporting, interrupts or traps from chipset errors, memory errors, sensor conditions, etc., up to the Linux Kernel level. All of these conditions shall be passed to the Linux Kernel level for handling. The delivered BIOS shall not attempt to respond these conditions directly other than a request to reboot or reset the node.

#### **3.2.5.12.6 BIOS Security (MR)**

The principal function of the provided BIOS shall be to perform node hardware initialization, power-on self-testing and turn over operation to the OS boot loader. The BIOS shall not perform any kind of extraneous automated or uncontrolled I/O to disks, networks, or other devices beyond that required to read or write BIOS images, CMOS parameters, and device or error registers as required for booting, BIOS configuration, power-on testing, and hardware diagnostics and configuration. Under no circumstances shall the BIOS itself be allowed to directly capture or write user data to any location. If the BIOS has such capabilities but is configured to disable them, then a formal testing process and the results of such tests shall be provided and conclusively demonstrate such features are in fact disabled.

### **3.2.5.12.7 Plans and Process for Needed BIOS Updates (MR)**

A written plan shall be submitted with offeror proposal(s) outlining Offeror's plans and process to provide BIOS updates to address problems or deficiencies in the areas of functionality, performance, and security. The plan shall outline a process which Offeror shall follow to identify, prioritize and implement BIOS updates in general and for addressing any specific Tri-Lab and/or DOE security related issues and concerns that may be raised. The Offeror's plan shall be finalized after subcontract award and included as an early deliverable the Statement of Work for the subcontract.

### **3.2.5.12.8 Failsafe/Fallback Boot Image (TR-1)**

The BIOS may employ a failsafe/fallback booting capability in case of errors in the default BIOS image, during the default BIOS boot process, or during a BIOS flash operation. The node may reach a minimum state allowing for the diagnosis of errors. This capability may not require the use of any external writeable media (e.g., floppy, flash disk, etc.) and preferably no external media at all.

### **3.2.5.12.9 Failsafe CMOS Parameters (TR-1)**

The BIOS may have failsafe default CMOS parameter settings so that the serial console interface will function if the CMOS values are unintentionally reset to the default values (e.g., the CMOS battery fails). The failsafe parameters may be modifiable in accordance with 3.2.5.12.13 below, if provided.

### **3.2.5.12.10 Serial Console after Failsafe/Fallback Booting (TR-1)**

After a failsafe/fallback BIOS condition (section 3.2.5.12.8) the BIOS would enable at a minimum a serial console which can operate with the remote management capabilities as specified in section 3.2.7 to allow remote access and management of the node. This capability may not require the use of any external media of any kind (CD-ROM, DVD, floppy, flash device, etc.).

### **3.2.5.12.11 BIOS Upgrade and Restore (TR-1)**

Offeror may provide a hardware and software solution that allows a Linux command line utility or utilities to update (flash) or restore the BIOS image(s) in the flash BIOS chip and to verify the flash image(s). This mechanism may not require booting an alternative operating system (e.g., DOS).

### **3.2.5.12.12 CMOS Parameter Manipulation (TR-1)**

The Offeror will provide a mechanism to read (get) and write (set) individual CMOS parameters from the Linux command line. In particular, CMOS parameters such as boot order may be modifiable from Linux command line. In addition, there may be accessible CMOS parameters that can disable any power management or processor throttling features of the node. The Linux command line utility shall also allow reading all CMOS parameters (backup) and writing all CMOS parameters (restore). A Linux based method to verify consistency of all CMOS parameters (excluding node or time dependent ones) shall be fully documented. Reading or writing CMOS parameters shall not require booting an alternative operating system or interacting with BIOS menus or a BIOS command line interface. Such a capability may function from within Linux command scripts and also allow the use with the Linux Expect utility and not require navigation of BIOS menus.

### **3.2.5.12.13 BIOS Command Line Interface (TR-2)**

A Linux shell command-line interface may be provided to interact with the delivered BIOS. It provides access to any other BIOS functionality above and beyond that described in section 3.2.5.12.12 and may be integrated with the tools associated with CMOS manipulation if both are provided. This capability may function from within Linux shell scripts and also may allow the use with the Linux Expect utility and may not require navigation of BIOS menus.

### **3.2.5.12.14 Serial Console over Ethernet Support (TR-2)**

The BIOS may directly provide a remote serial console capability over an Ethernet management network or over the IBA interconnect instead of over an actual serial port. This should support the remote interactive management features of section 3.2.7.

### **3.2.5.12.15 Power-On Self Test (TR-2)**

As a configurable option, the POST may be comprehensive, detect hardware failures and identify the failing FRU during the power up boot sequence. All POST failures should be reported to the serial console and through the remote management solution in section 3.3.12.

### **3.2.5.12.16 BIOS Security Verification (MR)**

The BIOS shall be provided with a vendor certification document indicating that the BIOS strictly meet the conditions specified in section 3.2.5.12.6. This certification document may become part of a DOE Security Plan for classified operations of these systems at the Tri-Laboratory site. Future updates of the BIOS shall be accompanied by a new certification appropriate for that specific BIOS version.

### **3.2.5.13 Programmable LED(s) (TR-3)**

Offeror will provide nodes with either programmable LED(s) or a programmable front Alpha Numeric panel for run-time diagnostics. Access to these may be made available to the BIOS as well as Linux.

## **3.2.6 IBA Interconnect (MR)**

The SU shall be built with an IBA compatible interconnect (<http://www.InfiniBandta.org/specs> version 1.2.1 with errata publications, or later). A network is IBA compatible if it is capable of running the RHEL software stack, section 3.3.1, 3.3.2) with a port of the HCA drivers and user level protocols (ULP's).

### **3.2.6.1 Node HCA Functionality (TR-1)**

The SU IBA network will have at least one IBA 4x QDR (or faster) with PCIe2 8x (or faster) HCA in every node. The HCA will not have local memory on the card (mem-free).

### **3.2.6.2 Node Bandwidth, Latency and Throughput (TR-1)**

The SU network will have at least one IBA 4x QDR (or faster) with PCIe2 8x (or faster) HCA. The SU IBA network shall deliver at least 90% of peak IBA network bandwidth,

both unidirectional and bidirectional, when exchanging data between any two nodes in the SU. For a 4x QDR network, this is at least 3.6 GB/s unidirectional (e.g. PingPong) and 7.2 GB/s bidirectional (e.g. SendRecv). The SU IBA network shall deliver an MPI ping/pong latency (round trip divided by two) of no more than 1.7  $\mu$ s as measured between any and all two MPI task pairs in the SU, each with one MPI task per core. The SU IBA network shall deliver an aggregate processing rate of at least  $1.7 \times 10^6$  messages per second per core utilizing an application with one MPI task per core on each node, between any two nodes in the SU. Subcontractor will report below the compute node delivered MPI bandwidth, latency and messaging throughput benchmarks over IBA 4x QDR (or faster) HCA attached to PCIe2 (or faster) buss between two nodes utilizing 1 MPI task/node, 1 MPI task/socket and 1 MPI task/core on each node.

The Tri-Laboratory suggests the “perfest”, “presta”, and “osu\_mbw\_mr” benchmarks from the following sites:

Perftest is available on the OpenFabrics website: <http://www.openfabrics.org>  
 Presta is available at: [https://asc.llnl.gov/computing\\_resources/purple/archive/benchmarks/presta/](https://asc.llnl.gov/computing_resources/purple/archive/benchmarks/presta/)  
 Osu\_mbw\_mr is available at: <http://mvapich.cse.ohio-state.edu/benchmarks/>

### 3.2.6.3 Fully Functional IBA Topology (TR-1)

The Offeror will propose a fully functional IBA 4x QDR (or faster) interconnect with all required hardware (e.g., switches, cables, connectors, and HCA) and software stack (section 3.3.1, 3.3.2). The delivered IBA network will be capable of running using RHEL 6.1 (TOSS 2.0) or later. The delivered IBA network hardware and software will be capable of driving copper or optical cables through the QSFP interface. If necessary, Offeror will assist the Tri-Lab, OpenFabrics, and RedHat community on porting the IBA software to the SU interconnects. For SU sizes up to 16 SU’s, the SU network will have at least one IBA 4x QDR (or faster) interconnect fabric with fat-tree topology and *full, non-blocking bisection bandwidth* that allows all nodes in the SU to communicate with every other node in the entire SU, possibly through spine switches. The SU network will have as many ports available for SU nodes as available for connection to spine switches in a *full, non-blocking bisection bandwidth* configuration. Larger switches (324-port or larger) are desired as the spine switches, but it may be beneficial from a cost standpoint to also use smaller (36-port or larger single switch ASIC) switches in the rack or as leaf switches.

### 3.2.6.4 IBA Cabling Pattern (TR-1)

The Offeror will connect compute nodes to the InfiniBand switches such that the successively numbered compute nodes can communicate through a minimum number of IBA ASIC switch chips (through a minimum number of hops). Cabling layout will be discussed with each deployment site and mutually agreed to prior to system build.

### 3.2.6.5 IBA Interconnect Reliability (TR-1)

In order to have 8 or 16 SU cluster aggregations that can support ASC Program production workloads, the Offeror will propose an IBA interconnect infrastructure that is reliable in the sense that it meets or exceeds the following:

- \* HCA failure requiring replacement is less than one per two months per 8SU cluster.

This corresponds to an HCA MTBF of over 1.558 million hours.

- \* HCA failures due to catastrophic state or reset or becomes non-responsive less than one per two months per 8SU cluster.
- \* Switch FRU failure rate less than one per two month per 8SU cluster. This corresponds to a switch FRU MTBF of over 473,000 hours.
- \* Less than one link loss per year per 8SU cluster. This corresponds to a link loss of over 7.4 million hours.
- \* Links drops back to below rated width or speed (for example 4x to 1x or QDR to DDR) no more than once per 1,024 restarts of Open SM per 8SU cluster.
- \* Cable Bit Error Rate (BER) better than  $1 \times 10^{-18}$
- \* IBA HCA PCI errors shall be 1 in 4 months and shall be reported as an error of the PCI bus rather than a generic error of the HCA.

### **3.2.6.6 Multi-SU Spine Switches (TR-1)**

The network for up to 4 SU clusters will be capable of directly connecting the SUs together without additional switches. The SU network will be capable of expanding to a cluster with at least up to 16 SU with the addition of Offeror supplied spine switches and cables.

### **3.2.7 Remote Manageability (TR-1)**

All nodes will be 100% remotely manageable, and all routine administration tasks automatable in a manner that scales up to a cluster aggregation of eight SU. In particular, all service operations under any circumstance on any node must be accomplished without the attachment of keyboard, video, monitor, and mouse (KVM). The Tri-Laboratory community intends to use the open source tools PowerMan and ConMan for remote management, and therefore the Offeror will propose hardware and software that is reasonably compatible with this software environment and provide any software components needed to integrate the proposed hardware with these tools. Areas of particular concern include remote console, remote power control, system monitoring, and system BIOS.

The Offeror will fully describe all remote manageability features, protocols, APIs, utilities and management of all node types bid. Any available manuals (or URLs pointing to those manuals) describing node management procedures for each node type will be provided with the proposal.

All remote management protocols, including power control, console redirection, system monitoring, and management network functions must be simultaneously available. Access to all system functions within the SU, must be made available at the Linux level so as to enable complete system health verification.

Offeror will propose a remote management solution that is based on section 3.2.7.1 or section 3.2.7.2, but not both. The Tri-Laboratory community has a strong preference for reliable, complete solutions based on section 3.2.7.2.

#### **3.2.7.1 Traditional Remote Management Solution (TR-1)**

IPMI 2.0 is preferred as the remote management solution. A fully functional traditional remote management solution (TRMS) may be proposed as a backup.

### **3.2.7.1.1 Remote Console (TR-3)**

The TRMS will interface to the console port of on every node in the SU. The TRMS console interfaces will be aggregated in Terminal servers on the Management Ethernet. The TRMS will interface to the power plug on every node of the SU. The TRMS power interfaces will be aggregated in a power control device on the Management Ethernet.

### **3.2.7.1.2 Remote Node Power Control over Management LAN (TR-3)**

Remote access to power control device over the management LAN will be accomplished through a command line interface that is can easily be scripted with the Linux Expect utility. The power control device will be capable of turning each node's power off and turning each node's power on and querying the power state of the node. The power control infrastructure will be able to reliably power up/down all nodes in the SU simultaneously. Reliable here means that 1,000,000 power state change commands will complete with at most one failure to actually change the power state of the target nodes.

### **3.2.7.2 IPMI and BMC Remote Management Solution (TR-1)**

Node remote management will be accomplished with IPMI 2.0 and a baseboard management controller (BMC). In the event that the BMC is not integrated to the base motherboard, the BMC daughter card (or equivalent) will be proposed. The Offeror will provide a fully compliant IPMI 2.0 implementation that will operate with FreeIPMI (<http://www.gnu.org/software/freeipmi>) including satisfying all IPMI specification mandatory requirements. All security relevant features in the IPMI specification must be supported and configurable. All IPMI functions will be utilized from Linux and there should be no requirements for any DOS based utilities. Offeror will provide a Linux command line utility or utilities that allow upgrade and verification of BMC firmware and BMC configuration values. The command line utility will allow reading of necessary BMC configuration parameters and writing of necessary BMC configuration parameters. Linux command line utilities for firmware upgrades must be able to work in-band. An out-of-band only firmware upgrade solution is not acceptable. BMC configuration will be based on FreeIPMI bmc-config utility (<http://www.gnu.org/software/freeipmi>). All security-relevant fields such as usernames, passwords, keys, user access, channel access, authentication, and enabling/disabling of features shall be configurable. Although it is not sufficient to ensure IPMI 2.0 compliance, it is highly recommended that the subcontractor verify that their systems pass at least the FreeIPMI IPMI compliance tests described in <http://www.gnu.org/software/freeipmi/freeipmi-testing.txt>. Offeror's proposed solution will not require OEM IPMI extensions for setup, monitoring or remote manageability. If IMPI OEM extensions are required, the subcontractor shall provide documentation on the extensions explaining additional commands, IPMI error codes, device error codes, sensors, system event log (SEL) events, sensor data repository (SDR) records, field replaceable unit (FRU) records, etc. so that they may be added into FreeIPMI. The documentation shall be detailed enough so that LLNS can understand the OEM extensions fully. Offeror will publicly release documentation on any OEM extensions (see <http://www.gnu.org/software/freeipmi/freeipmi-oem-documentation-requirements.txt> for OEM documentation requirements).

The IPMI solution will allow the following requirements below to be met concurrently over the SU management LAN.

#### **3.2.7.2.1 ConMan Access to Console via Serial over LAN (TR-1)**

ConMan will access all node consoles simultaneously via the IPMI 2.0 Serial Over LAN (SOL) for serial console access. The SOL session will be encrypted using AES-CBC-128 as defined in IPMI 2.0. The SOL implementation will meet requirements for serial console listed in Sections 3.2.7.3.1 through 3.2.7.3.3.

#### **3.2.7.2.2 LAN PowerMan Access (TR-1)**

PowerMan will access the BMC power management features on every node in the SU simultaneously via the FreeIPMI ipmipower tool. The BMC power management features will be capable of turning each node's power off and turning each node's power on. The BMC based power control infrastructure will be able to reliably power up/down all nodes in the SU simultaneously. Reliable here means that 1,000,000 power state change commands will complete with at most one failure to actually change the power state of the target nodes.

#### **3.2.7.2.3 LAN Management Access (TR-1)**

All other node management functions will be accomplished via a remote mechanism to every node in the SU simultaneously. The remote node management mechanism implementation will never allow message buffer overflow or data corruption conditions.

#### **3.2.7.2.4 Traditional Remote Management Backup Plan (TR-2)**

Offeror will propose a traditional remote management backup plan meeting the requirements in section 3.2.7.1 as a back up plan should the IPMI 2.0 based solution prove to be unworkable and/or unreliable.

#### **3.2.7.2.5 Additional IPMI Security Requirements (TR-1)**

Due to security policies in place at the Tri-Laboratories, the Offeror will provide several additional IPMI features not considered mandatory in the IPMI specification so that security policies can be met. The additional security features will be provided via IPMI commands and sensor events that are capable of being executed and read via FreeIPMI. IPMI commands and sensor events will be available to be published in the GPL software released by the Tri-Laboratories. Binary or web based tools that supply the features are not acceptable.

##### **3.2.7.2.5.1 Bad Password Threshold (TR-1)**

The Subcontractor will support "Bad Password Threshold", as defined by IPMI 2.0 Addenda and Errata E443 (See [http://download.intel.com/design/servers/ipmi/IPMI2\\_0E4\\_Markup\\_061209.pdf](http://download.intel.com/design/servers/ipmi/IPMI2_0E4_Markup_061209.pdf)).

This feature is listed as optional in the IPMI 2.0 Addenda and Errata, however it is considered a requirement for this procurement. When a user has surpassed the threshold, an appropriate "Session Audit" system event will be generated as defined by IPMI Addenda and Errata E443 and will be available for reading via FreeIPMI's ipmi-sel or platform event filtering tools. All BMC configuration settings will be done with FreeIPMI's bmc-config.

#### **3.2.7.2.5.2 Bad Password Monitoring (TR-1)**

The Subcontractor will support bad username and bad password "Session Audit" as defined by IPMI 2.0 Addenda and Errata E423 (See [http://download.intel.com/design/servers/ipmi/IPMI2\\_0E4\\_Markup\\_061209.pdf](http://download.intel.com/design/servers/ipmi/IPMI2_0E4_Markup_061209.pdf)).

When an invalid username or password has been specified, an appropriate system event will be generated. It will be available for reading via FreeIPMI's ipmi-sel or platform event filtering tools.

#### **3.2.7.3 Remote Management Solution Requirements (TR-1)**

The following requirements apply to both the IPMI 2.0 (section 3.2.7.2) and TRMS (section 3.2.7.1) solutions.

##### **3.2.7.3.1 Serial Console Redirection (TR-1)**

All BIOS interactions will be through a serial console. There will be no system management operations on a node that require a graphics subsystem, KVM or DVDROM or floppy to be plugged in. In particular, the serial console will display POST messages including failure codes, operate even upon failure of CMOS battery and provide for a mechanism to remotely access the BIOS.

##### **3.2.7.3.2 Dedicated Serial Console Communications (TR-2)**

The serial console communication channel on every node will be available simultaneously for console logging and interactive use at all times. This is to ensure that all console output is logged and Linux Expect scripts that perform console or service processor actions do not interfere with each other or with console logging. All console output must be available for logging at all times with no dropped or corrupted data. The serial console will operate after node crash or hang. The serial console will operate during a network boot.

##### **3.2.7.3.3 Serial Console Efficiency (TR-1)**

The serial console communication channel will support a baud rate of 115,200 or greater. If an IPMI 2.0 solution is offered, the BMC will transfer AES-128 encrypted character data at a rate equivalent to a traditional 115200 baud serial console.

##### **3.2.7.3.4 Flow Control (TR-1)**

Flow control will **not** be required for serial console communication.

##### **3.2.7.3.5 Peripheral Device Firmware (TR-2)**

Offeror will provide Linux utility or utilities for saving, restoring, verifying (including printing version number) firmware for any peripheral devices supplied.

##### **3.2.7.3.6 Remote Network Boot Mechanism (TR-1)**

The node BIOS will support booting an executable image over the management Ethernet utilizing PXE, BOOTP or DHCP with the vendor BIOS. Console data and power management functionality must be available during network boot process.

### **3.2.7.3.7 Serial Break (TR-1)**

The serial console subsystem shall be capable of transmitting and reliably delivering a “serial break” from a remote management station connected to the serial console solution through to the Linux kernel on each node of the SU. This functionality shall provide system administrators the ability to extract debug information from crashed nodes using kernel SysRq hooks.

### **3.2.7.4 Remote Management Security Requirements (TR-1)**

Additional features, such as ssh or web servers, are common in remote management solutions featuring IPMI, BMC, LOM, Out of band management, etc. These additional features open potential security issues, such as open ports. If additional features such as these are available, the offeror will provide a means to enable or disable these features.

The configuration shall be offered via a solution appropriate given the remote management solution provided by the vendor. For example, for an IPMI solution, a set of IPMI OEM commands to configure the current settings shall be made available.

### **3.2.7.5 GPU Node Requirements (MOR)**

As an MOR for LANL, the following requirements are specific to the GPU accelerated nodes and supersede the general node requirements above. The GPU nodes would replace each of the compute nodes with a GPU enabled equivalent node. These SUs will be used as hybrid compute nodes. We prefer a minimum of two GPUs per node (more is better). The vendor shall demonstrate that the proposed solution will work for this use case.

### **3.2.7.6 GPU Node General Requirement (MOR)**

As an MOR for LANL, the GPU-enhanced SUs shall utilize the same processors, BIOS, memory and IB interconnect as the Compute Node. There shall be sufficient PCIe2 (or better) x16 slots to accommodate the number of GPUs specified, plus the IB. All slots shall be capable of operating at full bandwidth simultaneously.

### **3.2.7.7 GPU Node Architecture (MOR)**

As an MOR for LANL, each GPU shall have at least 2GB of ECC capable memory (more GPU memory is desirable). The GPU shall have fast double precision performance (more aggregate GPU double precision performance is desirable). Manageability features on the GPU are preferred.

### **3.2.7.8 GPU Node Memory Requirement (MOR)**

As an MOR for LANL, the CPU memory requirement (**not** GPU memory requirement) for the GPU-enhanced nodes shall be at least 2GB times the number of CPU cores. A memory option to add more CPU memory is desired.

## **3.2.8 Gateway Node Requirements (TR-1)**

The following Requirements are specific to the gateway nodes and supersede the general node requirements above.

### **3.2.8.1 Gateway Node Count (TR-1)**

The Offeror will configure the SU with six (6) Gateway nodes for SU peak up to 100 teraFLOP/s and twelve (12) Gateway nodes for SU peak above 100 teraFLOP/s.

### **3.2.8.2 Gateway Node Configuration (TR-1)**

The SU gateway nodes will be used for file system and other IP based connectivity between the compute nodes and the global file system and other IP based communications infrastructure. For Lustre, the gateway will run the “LNet Router code,” that routes LNet/Verbs/IBA to LNet/IBA. For PanFS, the gateway will route IP packets between IP/IBA to IP/10 Gb/s Ethernet using Quagga/Zebra Linux OSPF routing software. Offeror will configure the SU with a minimum number of gateway nodes to achieve a delivered gateway bandwidth throughput of 0.0004 B:F, where B is the aggregate number of Bytes/s that the gateways can route IP packets between the IBA and 10 Gb/s Ethernet networks. For a 50 teraFLOP/s SU this requirement translates into an aggregate gateway delivered IP routing bandwidth of 20 GB/s. The Offeror’s gateway should carefully balance the IBA HCA and 10 Gb/s Ethernet network *delivered* IP bandwidths. Gateway nodes may also include at least one 1 Gb/s 1000Base-TX Ethernet interface in addition to any management Ethernet. This interface will be used to route NFS traffic between compute nodes and NFS file servers on the 1-10 Gb/s Ethernet infrastructure.

### **3.2.8.3 Gateway Node I/O Configuration (TR-1)**

The SU gateway nodes chipset will be configured with sufficient PCIe2 8x (or faster) busses and sufficient slots to drive both the SU internal IBA HCA and the SAN network interfaces to either 4.0 GB/s (IB) or four 10 GbE. The gateway node must be capable of driving 2 IB QDR (or faster) interfaces or one IB QDR (or faster) interface and four 10GbE interfaces at the same time at least 90% of peak. The SU gateway node PCIe2 (or faster) infrastructure will be fully compatible with the proposed IBA HCA and network cards.

### **3.2.8.4 Gateway Node QDR IB Card (TR-1)**

The SU gateway nodes will be configured with one (or two) PCIe2 8x (or faster) 4x QDR InfiniBand card(s). Offeror will provide and support an open source IB driver for Linux 2.6 kernels. The IB card(s) will natively support standard IB protocols such as: IB Verbs, IPoIB, SRP, and iSER. All network interfaces and device drivers will support 9KB jumbo frame or greater MTU operation.

### **3.2.8.5 Gateway Node 10Gb Ethernet Card (TR-1)**

The SU gateway nodes will be configured with four 10 Gb Ethernet ports with SR optics, capable of delivering an aggregate 40 Gb/s peak performance and at least 90% of peak. Offeror will provide and support an open source Ethernet driver for Linux 2.6 kernels.

### **3.2.8.6 Gateway Node Delivered Performance (TR-2)**

The SU gateway nodes will be configured to support a minimum of 4 GB/s IBA to IBA delivered IB routing bi-directional bandwidth (counting both directions). Offeror will provide fully documented benchmark data demonstrating the minimum performance utilizing the NTTCP benchmark with the Offeror’s response.

### **3.2.9 Login/Service/Master Node Requirements**

The following Requirements are specific to the Login/Service/Master (LSM) nodes and supersede the general node requirements, above. LSM nodes will be used for management functions as well as user access (e.g., Login, application development and job launch).

#### **3.2.9.1 LSM Node Count (TR-1)**

The Offeror will configure the SU with one LSM node for SU up to 100 teraFLOP/s and two LSM nodes for SU peak above 100 teraFLOP/s. Sites may wish to negotiate a specific number of LSM nodes on a per cluster basis, Offeror will support this configuration.

#### **3.2.9.2 LSM Node I/O Configuration (TR-1)**

The SU LSM nodes chipset will be configured with sufficient PCIe2 (or faster), busses and sufficient PCIe2 slots to drive the IBA HCA, and 1 Gb/s Ethernet cards at full line rate. The SU LSM node PCIe2 (or faster) infrastructure will be fully compatible with the proposed IBA HCA, and 1 Gb/s Ethernet cards.

#### **3.2.9.3 LSM Node Ethernet Configuration (TR-1)**

The SU LSM nodes will be configured with one (1) PCIe2 (or faster) 4x QDR InfiniBand card for access to the site IB infrastructure. The SU LSM nodes will be configured with dual 10 Gb/s multi-mode or single-mode (depending on site preference) fiber ports supported by a PCIe2 or better bus for access to the site 10 Gb/s Ethernet infrastructure. These 10 Gb/s Ethernet and 40 Gb/s IB ports will be in addition to any ports required for management functions. Offeror will provide and support open source 10 Gb/s Ethernet and 4x QDR IB drivers for Linux 2.6.31 and later kernels. All network interfaces and device drivers will support 9KB jumbo frame or greater MTU operation.

#### **3.2.9.4 LSM Node Accessory Configuration (TR-2)**

The SU LSM nodes will be configured with one (1) read only (not read/write) 4x DVD-ROM bootable device. The SU LSM nodes will be configured with at least one (1) 1.5 TB (or larger) SATA disk. This disks should configured in a reliable manner (at least RAID 1) and be hot swappable as well as directly accessible from the exterior of the LSM node.

### **3.2.10 Remote Partition Service Node Requirements**

The following Requirements are specific to the Remote Partition (RPS) nodes and supersede the general node requirements, above. RPS nodes will be used as a remote disk device for the compute and gateway.

#### **3.2.10.1 RPS Node Count (TR-1)**

The Offeror will configure the SU with one (1) RPS node for SU up to 100 teraFLOP/s and two (2) RPS nodes for SU peak above 100 teraFLOP/s.

#### **3.2.10.2 RPS Node I/O Configuration (TR-1)**

The SU RPS nodes chipset will be configured with sufficient PCIe2 (or faster), busses and sufficient slots to drive the IBA HCAs, Ethernet card at full line rate. The SU RPS

node PCIe2 and HyperTransport™ infrastructure will be fully compatible with the proposed IBA HCA, and Ethernet cards.

### **3.2.10.3 RPS Node Ethernet Configuration (TR-1)**

The SU RPS nodes will be configured with one (1) PCIe2 (or faster) 4x QDR InfiniBand card for access to the site IB infrastructure. The SU RPS nodes will be configured with dual 10 Gb/s Ethernet multi-mode SR fiber ports supported by a PCIe2 or better bus for access to the site 10 Gb/s Ethernet infrastructure and access to the management Ethernet infrastructure. These 10 Gb/s Ethernet ports and 40 Gb/s IB ports will be in addition to any ports required for RPS node management functions. Offeror will provide and support open source 10 Gb/s Ethernet and 4x QDR IB drivers for Linux 2.6.31 and later kernels. All network interfaces and device drivers will support 9KB jumbo frame or greater MTU operation.

### **3.2.10.4 RPS Node RAID Configuration (TR-1)**

The SU RPS nodes will be configured with at least one (1) highly reliable hardware RAID configuration utilizing at least four (4) active 15K RPM SAS disks with aggregate capacity of at least 500 GB, plus an identical hot spare disk. The RAID controller may be capable of at least RAID5, RAID6, and RAID10. The RAID arrays will deliver at least 1 GB/s (best case, using outer cylinders) and 500 MB/s (minimum, using inner cylinders) aggregate sustained read/write bandwidth, and at least 90% of that performance should be obtainable from the Linux EXT4 file system mounted on each partition. The RAID arrays will deliver an average seek time of better than 4ms and average latency better than 2.5 ms. The individual disks will feature nonrecoverable read errors of 1 sector per  $10^{16}$  bits or better and MTBF rating of at least 1.2 million hours.

### **3.2.11 SU Management Ethernet (TR-1)**

The Offeror will provide a management Ethernet 1000Base-TX (copper) for the SU. The management Ethernet infrastructure will provide access to every node. In the case of failure in the IBA interconnect, the management network can be used to boot the entire system. The management Ethernet will be aggregated with high quality, high reliability Ethernet switches with full bandwidth backplanes and provide a single 1000Base-TX (copper) Ethernet uplink. The management Ethernet cables will be bundled within the rack in such a way as to not kink the cables, nor place strain on the Ethernet connectors. All management Ethernet connectors will have a snug fit when inserted in the management Ethernet port on the nodes and switches. The management Ethernet cables will meet or exceed Cat 5E specifications for cable and connectors. Cable quality references can be found at: (<http://www.integrityscs.com/index.htm>) and ([http://www.panduitnecg.com:80/whatsnew/integrity\\_white\\_paper.asp](http://www.panduitnecg.com:80/whatsnew/integrity_white_paper.asp)).

Offeror will provide CAT6 or equivalent management cables. A suggested source of this quality cable is Panduit corporations Powersum+ tangle free patch cords, Part# UTPCI10BL for a 10' cable. The URL for this product is: ([http://www.panduitnecg.com/solutions/copper\\_category\\_5e\\_5\\_3.asp](http://www.panduitnecg.com/solutions/copper_category_5e_5_3.asp)).

Management Ethernet reliability is specified in Section 4.1.

### **3.2.12 SU Racks and Packaging (TR-1)**

The Offeror will place the TLCC2 SU nodes, global disks, RAID controllers, IBA infrastructure, management Ethernet, in standard computer racks with ample room for

cable management of InfiniBand cables, CAT5e or CAT6 Ethernet cables and console serial cables and power cables. There will not be any console display, keyboard or mouse (KVM) equipment in the racks, except in the rack containing the LSM node. The LSM node, in each SU, will be connected to a single rack mount 1U keyboard, monitor and mouse. For a 324 or 648 port IB switch the rack should be a 30" wide rack.

#### **3.2.12.1 SU Design Optimization (TR-2)**

Offeror's SU design will be optimized to minimize the overall footprint of 2, 4, 6, 8 and 16 SU aggregations within the other constraints in section 3.2.12.

#### **3.2.12.2 Rack Height and Weight (TR-1)**

SU will not be taller than 84" high (48U) and not place an average weight load of more than 250 lbs/ft<sup>2</sup> over the entire footprint of the SU, including hot and cold aisles. If Offeror proposes a rack configuration that weighs more 250 lbs/ft<sup>2</sup> over the footprint of the rack, then Offeror will indicate how this weight can be redistributed over more area to achieve a load less than 250 lbs/ft<sup>2</sup>. Note also the site delivery and facility restrictions in Section 0.

#### **3.2.12.3 Rack Structural Integrity (TR-2)**

The provided racks will be of high structural quality. In particular, rack frames will be of sufficient strength and rigidity that the racks will not flex nor twist under the external load of a human being pushing at eye level on the rack from any of the four corners or sides. Additional reinforcement will be added as necessary to maintain rack structural integrity.

As a seismic event precaution, upon SU delivery Offeror will bolt racks in each row together with at least four 3/4" lag bolts, or better, for end racks and eight 3/4" lag bolts, or better, for racks sited touching two other racks (one bolt on each side corner touching another rack). During SU assembly at Offeror's facility, racks should have holes for inter rack bolting drilled prior to the placement of ANY equipment in them. SU's sited next to existing equipment (e.g., prior SU deliveries) in the same row need not be bolted together. The rack base will have wheels, leveling feet and adequate structural integrity to allow the rack to be bolted internally through the computer floor to the concrete sub-floor where required. The rack base must also allow adequate hole penetration for power and communication cables. The rack top must allow for overhead cable routing.

#### **3.2.12.4 Rack Air Flow and Cooling (TR-1)**

The racks will have sufficient airflow to adequately cool at full load the equipment mounted in the rack and racks installed at 600 ft. (LLNL and SNL California) and 7,500 ft. (LANL) or 5,400 ft. (SNL New Mexico) elevation with 30% humidity at up to 60° F (LLNL and SNL California) and 60° F (LANL or SNL New Mexico) air intake temperature. Where necessary, the rack bottom panel will be completely removed to improve airflow and allow sufficient room to run cabling out of the rack and under the floor.

#### **3.2.12.5 Rack Doors (TR-2)**

The rack, if provided with a front or rear door, will include a non-breakable, see through panel (such as a metal mesh or grid) and have sufficient perforations to maintain adequate

airflow throughout the cabinet while closed. The front and rear rack doors will be lockable. Where additional cooling is necessary (see section 3.2.12.8), liquid cooled doors are desired.

### **3.2.12.6 Rack Cable Management (TR-2)**

The racks will have sufficient room for all equipment and cables without impeding airflow through the rack. All cables within a rack will be supported so that the weight of the cable is not borne by the cable attachment mechanism. A rubber grommet or other protection will be placed around the rack bottom opening as necessary to protect the IBA and other cables from damage. In addition, cables will be attached to rack mounts installed in the rear and/or front of the cabinet for cable management. Cable management solution may not block access to active components in the rack. Rack cabling shall allow the removal of any FRU in the rack without having to significantly uncable or recable the entire rack.

### **3.2.12.7 Rack Color (TR-3)**

All racks will be black and covered with a fully powder coated style paint or other covering.

### **3.2.12.8 Rack Power and Cooling (TR-1)**

Overall power and cooling for the SU are TCO components for the Tri-Laboratory community. For racks with air cooling solutions that require all the cooling from air provided by the facility, each rack will not require more than 50 kW (LLNL), 24 kW (LANL or SNL) of power, and corresponding cooling, assuming front to back air cooling.

Separate LANL requirement is for a third rack configuration which will either a) use liquid cooling capable of removing 100% of the heat generated, or b) not require more than 10 KW of power, and corresponding cooling, assuming front to back air cooling. Liquid cooling is the preferred and anticipated solution. If the rack requires more than the above power envelopes, then Offeror will propose less dense solutions and/or alternative cooling apparatus that reduces the intake air-cooling load. Offeror will fully describe the liquid cooling apparatus and the implications for siting and facilities modifications (e.g., chilled water feeds, flow rates). Specific cooling solutions of interest include rack-mountable, liquid cooled doors.

### **3.2.12.9 Rack PDU (TR-1)**

Rack PDU for the SU will minimize the number of 208V circuit breakers required in wall panels at the Tri-Laboratory sites. One (1) would be ideal, but the per-circuit limit depends on the installation site. In addition, the amperage of the required circuit breakers should be calibrated so that the utilization is maximized, but below 80% of the rated load during normal operation with heavy workload of user applications running. If the equipment in the rack requires more power during power-up (so called surge power), the rack PDU shall not trip circuit breakers under normal power-up conditions. The sustained PDU load need not be calibrated to this surge power, but rather to the normal operating power with user applications running.

The rack loads should be connected to the rack PDU so that the connected load is equally balanced across each phase. The phase imbalance of the total rack load shall be no

greater than 5%.

The Rack PDU will have on-off switches or switch rated circuit breakers to allow system administrator to power down all components in a rack with switches or circuit breakers in the PDU.

### **3.2.13 Safety Standards and Testing (TR-1)**

Materials, supplies, and equipment furnished or used by the Offeror under this SOW shall meet nationally recognized safety standards or be tested by the Offeror in a manner demonstrating they are safe for use. All electrical equipment, components, conductors, and other electrical material shall be of a type that is listed, labeled, or tested by a Nationally Recognized Testing Laboratory (NRTL) in accordance with Title 29, Part 1910, *Occupational Safety and Health Standards*, of the Code of Federal Regulations (29 CFR 1910). The Offeror shall obtain prior written approval from the LLNS Contract Administrator before furnishing or using any materials, supplies, or equipment under this SOW that do not meet these requirements.

## **3.3 TLCC2 Software Requirements (TR-1)**

This section describes the software requirements beyond the Government Furnished Software (GFS) for the TLCC2. The software associated with building and installing the TLCC2 SUs is described in Section 2.3. Offeror will provide all source code as Open Source in the form of *buildable* source SRPMs with the provided software.

### **3.3.1 Minimum IBA Software Stack (MR)**

The Offeror shall provide and support a fully compliant IBA V1.2.1, or later with published errata (<http://www.InfiniBandta.org/specs>) Linux software stack for the TLCC2 SU. The Offeror's IBA software stack shall be fully functional, stable and scale on clusters comprised of aggregations of 1 to 16 SUs. The IBA software stack shall include at least the following components: HCA driver, core InfiniBand modules (subnet management agent, performance management agent, connection manager, subnet administration, general service interface), kernel and user space Verbs, IPoIB (Connected mode), SRP or iSER (initiator and target), MPI, OpenSM, network and host diagnostics (InfiniBand-diags) and firmware for HCA and IBA switches and Linux command line utilities for flashing the HCA firmware from the node it is attached to and the switch firmware LSM node over the management Ethernet.

#### **3.3.1.1 IBA Software Stack Compatibility (MR)**

The Offeror's supplied and supported IBA software stack shall be RHEL 6.x based. Extensions to the IBA software shall be compatible with the TOSS 2.0 or later kernel (2.6.32 or later).

The Offeror's IB software stack shall be deemed production quality by the Tri-Laboratory community if it successfully completes the Tri-Laboratory (pre-ship, post-ship and/or acceptance) workload test plan exit criteria on the proposed hardware at 1 SU scale.

Functionality beyond the current IBTA (InfiniBand Trade Association) specification shall maintain compatibility with that specification thus allowing for maximum

interoperability among IB hardware. In addition, any “proprietary” extensions shall have an open source (GPL/BSD license) or open API solution for their use.

### **3.3.1.2 Open Source IBA Software Stack (TR-1)**

The Offeror shall contribute all modifications to the IB software stack to the open source community (OpenFabrics Alliance and RedHat) throughout the lifetime of this procurement.

IBA diagnostics shall be accessible by open source tools such as those provided by the “InfiniBand-diags” open source package.

### **3.3.2 IBA Upper Layer Protocols (TR-2)**

The Offeror's provided and supported InfiniBand stack releases will also include the following Upper Layer Protocols (ULP):

SDP ([www.rdmaconsortium.org/home](http://www.rdmaconsortium.org/home))

user space DAPL, <http://www.datcollaborative.org/udapl.html>

IPoIB, [http://www.ietf.org/html.charters/OLD/ipoib-](http://www.ietf.org/html.charters/OLD/ipoib-charter.html)

[charter.htmlhttp://www.datcollaborative.org/kdapl.html](http://www.datcollaborative.org/kdapl.html)

SRP, <http://www.t11.org/t10/drafts/srp/srp-r16a.pdf>

iSCSI, <http://www.ietf.org/rfc/rfc3720.txt>

iSER, <http://www.rdmaconsortium.org/home>

NFS-RDMA, <http://www.ietf.org/rfc/rfc3010.txt>

IPoIB connected mode, <http://www.ietf.org/internet-drafts/draft-ietf-ipoib-connected-mode-00.txt>

These protocols will fully implement and conform to the above specifications.

### **3.3.2.1 TLCC2 IB HCA Error Reporting (TR-3)**

Hardware errors detected by the HCA, which are not the direct responsibility of the HCA (for example PCI errors), will be reported as such by the FW/Driver of the HCA. PCIe error reporting shall be enabled by the BIOS to help facilitate this.

### **3.3.2.2 TLCC2 IB Switch Firmware Update (TR-3)**

Offeror will provide open source command-line tools to flash switch firmware over the IB network.

### **3.3.3 TLCC2 Peripheral Device Drivers (TR-1)**

Offeror will provide Linux drivers for all peripheral devices supplied that function with the TOSS kernel. This additional or modified software must be provided as source or as buildable source RPMs with licensing terms which allow for the free redistribution of that source (BSD or GPL preferred). Offeror will specifically call out and fully disclose any proposed peripheral device drivers required with the proposed SU including version number and provide system administration or programmers documentation with the proposal.

### **3.3.4 GPU Node Software (MOR)**

The GPU enhanced nodes shall run the same basic CCE software stack as the rest of the system. The Offeror shall provide all required proprietary and optimized Linux drivers for

the GPU hardware, as well as any GPU vendor diagnostics. The provided drivers shall support all OpenCL functionality.

In addition for the GPU-Enhanced clusters, the LSM nodes shall be capable of accessing the appropriate libraries and compiling code for GPU execution.

### **3.3.5 RPS Node Software (TR-1)**

The RPS node will provide remote access to root and /swap file systems for compute and gateway nodes. As the compute and gateway nodes boot over the management or IBA network, they will establish connections to the RPS node and mount their root (including /tmp and /var/tmp) (EXT3 file system) and /swap (block device) partitions from the RPS node via SCSI Remote Protocol (SRP) or iSCSI with extensions for RDMA (iSER) target over IBA. Thus, the Offeror will supply the SRP or iSER target code for the RPS node. In addition, Offeror will supply the RAID5 device driver for the RPS node RAID5 controller that is compatible with the RHEL6 LVM layer. With this RPS node software configuration it will be possible to simultaneously service root and /swap partitions for all compute and gateway nodes.

### **3.3.6 Hardware Memory Uncorrectable Error Detection (TR-1)**

The SU nodes will include a hardware mechanism to detect memory uncorrectable errors. This hardware mechanism will be capable of sending a non-maskable interrupt (NMI) or machine check exception when an uncorrectable error occurs, so that the Linux operating system may take immediate action. When a memory uncorrectable error occurs, this hardware mechanism will provide sufficient information so that the Linux operating system may identify the affected failed or failing memory component FRU (i.e. the exact DIMM FRU identified by the label visible on the motherboard) and log it without requiring an atypical reboot or a manual procedure to recover the error from a system event log.

### **3.3.7 Hardware Memory Corrected Error Detection (TR-1)**

The SU nodes will include a hardware mechanism to detect and count memory corrected errors. The node memory controller hardware mechanism that detects and counts memory corrected errors may have low system overhead in that it will utilize less than 1 processor core cycle or memory bus transaction per million cycles or bus transactions when memory corrected error rates are less than one per minute. When a memory corrected error occurs, this hardware mechanism will keep track of sufficient information so that the Linux operating system may identify the affected memory component FRU(s) (i.e. the exact DIMM FRU identified by the label visible on the motherboard). These hardware counters need not be perfectly accurate but should be able to appropriately reflect the enormity of the detected problem even when confronted with very high rates of memory corrected errors. In other words, while some loss of events may be unavoidable, high rates of memory corrected errors should not cause substantial undercounts. Node memory controller will expose a publicly open and documented interface that will allow the Tri-Laboratory personnel to write a Linux command line utility that runs at the Linux shell prompt and can directly obtain the type of correction mechanism actually enabled (e.g. Simple ECC SECDED vs. Chipkill X4 vs. Chipkill X8), the memory corrected error counts per memory component FRU and also reset all the error counts to zero. In addition, node memory controller will expose defective memory modules that the chipkill or other

functionality compensates for to the Linux operating system down to the memory component FRU(s) level.

### **3.3.8 Hardware Memory Controller Capabilities & Configuration (TR-1)**

The node memory controller configuration capabilities and option settings that exist may be directly exposed through a publicly open interface to the Linux operating system. The term “directly exposed” specially precludes the interposition of the node's BIOS and/or SMBIOS tables and/or alternative hardware mechanisms other than the memory controller(s) itself (themselves) in between the state of the memory controller and the provided interface.

The provided memory controller and interface may be sufficiently well implemented and complete so that Tri-Laboratory personnel can implement a kernel driver or module that exports these settings and where appropriate the ability to manipulate them to user space for verification during the boot process via a Linux command line utility. These memory controller configuration capabilities and options exposed through this interface may include but are not limited to: actual memory error detection state (enabled/disabled), the type of correction mechanism actually enabled (e.g., simple ECC SECDED, chipkill x4 or chipkill x8), memory RAS and CAS timing setting, memory chip and memory bus speed, memory interleaving, memory mirroring, and whether the memory controller or system firmware does scrubbing of corrected memory errors.

Offeror may provide hardware memory controller capabilities and configuration interface documentation that is sufficiently well written and complete so that Tri-Laboratory personnel can actually query and program this hardware facility will be delivered with the SU.

### **3.3.9 Software Support for Memory Error Detection and Configuration (TR-1)**

The Offeror will modify the Error Detection and Correction (EDAC) code in the Linux kernel or loadable kernel module to support the chipsets proposed (Sections 3.2.5.7-9). See <http://sourceforge.net/projects/bluesmoke/> for more information on the Linux kernel's support for memory EDAC. The Offeror will work with the Tri-Laboratory community to integrate this code into the Tri-Laboratory Linux distributions. This EDAC software with Offeror supplied modifications will log all memory uncorrectable errors to the Linux kernel log facility. This modification will report the failed or failing memory component FRU (i.e., the exact DIMM FRU identified by the label visible on the motherboard will be indicated in the kernel panic message). This EDAC software with Offeror supplied modifications will panic the node if the memory subsystem generates an uncorrectable memory error that the operating system cannot recover from. This EDAC software with Offeror supplied modifications will provide an appropriate Linux command line utility and interface to the hardware memory controller to query hardware memory controller correctable error counts and reset those counters (Section 3.2.5.8). This EDAC software with Offeror supplied modifications will provide an appropriate Linux command line utility and interface to the hardware memory controller to directly query hardware memory controller configuration (Section 3.2.5.9). This EDAC software with Offeror supplied

modifications will provide an appropriate interface to the hardware diagnostics in Section 4.5.2.

If the design of the hardware memory controller is so architecturally different from existing memory controllers that its information cannot be represented using the existing EDAC data structures, then the Offeror will either provide an open source, functionally equivalent sysfs interface and open source modifications to the edac-utils user space package to work with this new interface or Offeror will work with the upstream bluesmoke/EDAC developers to rearchitect EDAC's data structures to accommodate the hardware memory controller's architecture.

### **3.3.10 Memory Diagnostics (TR-1)**

The Linux OS will interface to the SU node hardware memory error facility specified in Section 3.3.5 to log all correctable and uncorrectable memory errors on each memory FRU. If the operating system cannot recover from an uncorrectable memory error without impacting the computational job, the Linux kernel will report the failing memory FRU and then panic the node. The Offeror will provide a Linux command line utility that can scan the nodes and directly query the memory controller on each node to determine corrected and uncorrectable memory error counts and identify at the FRU level indicating the exact memory component FRU identified by the label visible on the motherboard where the failing or failed memory component is located and reset the counters. The Offeror will provide a Linux command line utility that can scan the nodes and directly query the memory controller on each node to determine the precise memory configuration of the memory subsystem on that node.

### **3.3.11 Linux Access to Motherboard Sensors (TR-1)**

All IPMI sensor data will be accessible both in-band and out-of-band through FreeIPMI. Offeror will provide any changes required to FreeIPMI. These changes will be provided by the Offeror for inclusion in the open source FreeIPMI project. All power supply, processor state, and sensors listed in the Sensor Type Codes table (Table 42-3 of the IPMI 2.0 Specification) will supply the sensor values corresponding to those given in that table. The Offeror may not provide their own sensor values and interpretations.

If a TRMS (section 3.2.7.1) is proposed, Offeror will provide at least a command line mechanism for sampling motherboard sensor values for those listed in section 3.2.7.1 from within the Linux operating system. Sensor types and values will be output in such a way that scripts can parse them. The LM-SENSORS package (<http://secure.netroedge.com/~lm78/>) is one solution used by the Tri-Laboratory community. If LM-SENSORS is proposed, Offeror will provide any needed kernel device drivers under open source license and a correctly calibrated sensors.conf file, including threshold values that adhere to manufacturers specifications, for all node types offered.

The motherboard hardware shall provide the following sensor data:

- Each and every fan within the node
- Temperature of every processor die
- All motherboard temperature sensors
- Voltage supply to each socket
- Processor state
- Power supply state

Temperature sensors shall be designed to be insensitive to manufacturing tolerances, e.g., CPU thermal diode readings shall utilize the dual-sourcing current or more accurate methodology. Regardless of the sensor solution provided, the subcontractor shall publicly release documentation on any OEM specific motherboard sensors so that the sensors can be interpreted correctly.

Sensor accuracy, precision, and physical meaning shall be stated for each sensor. An individual sensor shall be provided for each power supply and processor (or processor core) that exists in the system. A single sensor that represents multiple power supplies or processors/cores is not acceptable.

### **3.3.12 Remote Management Software (TR-2)**

The Offeror will provide remote management software, beyond that defined in Section 3.2.7, for the remote management of the TLCC2 Cluster. This may include utilities to capture and monitor BIOS, Linux Console and other node management I/O. It is preferred that any provided software be Open Source.

#### **3.3.12.1 TRMS Software (TR-1)**

Offeror will supply software that handles remote power control over LAN and interfaces with PowerMan software. Offeror will deliver software that handles remote console access and logging over LAN and interfaces with the ConMan software.

#### **3.3.12.2 IPMI and BMC Remote Management Software (TR-1)**

If the Offeror proposes IPMI and BMC remote management solution (section 3.2.7.2), then Offeror will supply FreeIPMI software as the software component of the IPMI and BMC remote management solution (section 3.2.7.2), if proposed.

#### **3.3.12.3 Linux Tool for BIOS Upgrade (TR-1)**

The node BIOS will be delivered with a Linux command line tool for BIOS upgrade.

If Offeror proposes a standard BIOS, the MTD kernel device driver (<http://www.linux-mtd.infradead.org>) is one solution to this requirement that is already used by the Tri-Laboratory community. MTD offers a UNIX character device interface to common flash memory technology devices. On the TLCC2 SUs, an MTD based solution would be combined with scripts that safely implement flash/verify functions for the Offeror's BIOS images. If this method is proposed, all modifications to the MTD device driver and scripts will be provided to the Tri-Laboratory community under open source license. If another mechanism to meet this requirement is proposed, then the proposed tool will be offered to the Tri-Laboratory under open source license. It is not acceptable to deliver a tool that only works in an operating system other than Linux (e.g., DOS or Windows).

### **3.3.13 System Diagnostics (TR-2)**

See Section 4 for the list of system monitoring and diagnostics required.

## **3.4 SU Hardware Evolution (TR-1)**

For SU technology components that Offeror rated with “medium” or “high” impact in section 3.1.3, Offeror will list how those changes to the proposed solution will change the offering relative to the hardware requirements in section 3.2. Offeror’s response should indicate changes to meeting each requirement in section 3.2.X, with designation 3.4.X.

## **3.5 SU Software Evolution (TR-1)**

For SU technology components that Offeror rated with “medium” or “high” impact in section 3.1.3, Offeror will list how those changes to the proposed solution will change the offering relative to the software requirements in section 3.3. In addition, any software component offering that changes over the 3QCY11 through 1QCY12 time frame that has “medium” or “high” impact should be listed as well. Offeror’s response should indicate changes to meeting each requirement in section 3.3.X, with section heading 3.5.X.

**End of Section 3**

## 4 Reliability, Availability, Serviceability (RAS) and Maintenance

The TLCC2 SUs, aggregated into clusters, are intended for classified production usage at the Tri-Laboratory sites. The Tri-Laboratories therefore requires that the SUs have highly effective, scalable RAS features and prompt hardware and software maintenance. In addition, Offeror shall propose end-to-end support for the proposed IBA interconnect hardware and software.

For hardware maintenance, the strategy is that Tri-Laboratory personnel will provide on-site, on-call 24x7 hardware failure response. The Tri-Laboratory envisions that these hardware technicians and system administrators will be trained by the Offeror to perform on-site service on the provided hardware. For easily diagnosable node problems, Tri-Laboratory personnel will perform repair actions in-situ by replacing FRUs. For harder to diagnose problems, Tri-Laboratory personnel will swap out the failing node(s) with on-site hot spare node(s) and perform diagnosis and repair actions in the separate Hot-Spare Cluster (HSC). Failing FRUs or nodes (except for writable media) will be returned to the Offeror for replacement. Hard Disks FRUs and writeable media (e.g., EEPROM) from other FRUs will be destroyed by each Laboratory according to DOE/NNSA computer security orders. Thus, each Tri-Laboratory site requires an on-site parts cache of all FRUs and a small cluster of fully functional hot-spare nodes of each node type. The Offeror will work with the Tri-Laboratory community to diagnose hardware problems (either remotely or on-site, as appropriate). On occasions, when systematic problems with the cluster are found, the Offeror's personnel will augment Tri-Laboratory personnel in diagnosing the problem and performing repair actions.

In order for the large number TLCC2 SUs to fulfill the mission of providing the capacity resource for the ASC Program and SSP, they must be highly stable and reliable from both a hardware and software perspective. The number of failing components per unit time (weekly) should be kept to a minimum. SU components should be fully tested and burned in before delivery (initially and as FRU or hot-spare node replacement). In addition, in order to minimize the impact of failing parts, the Tri-Laboratory community must have the ability to quickly diagnose problems and perform repair actions. A comprehensive set of diagnostics that are actually capable of exposing and diagnosing problems are required. It has been our experience that this is a difficult but achievable goal, and the Offeror will need to specifically apply sufficient resources to accomplish it.

For software, the strategy is similar to the hardware strategy in that Tri-Laboratory personnel will perform the Level 1 and Level 2 software support functions. Specifically, Tri-Laboratory personnel will diagnose software bugs to determine the failing component. The problem will be handed off to the appropriate Tri-Laboratory organization for resolution. For Tri-Laboratory supplied system tools, Tri-Laboratory personnel will fix the bugs. For Offeror-supplied system tools, the Offeror will need to supply problem resolution. For the Linux kernel and associated utilities, the Tri-Laboratory community intends to separately subcontract with Red Hat for Enterprise level support. For file system related SW problems, the Tri-Laboratory community intends to separately subcontract with Cluster File Systems, Inc for Lustre support and with Panasas for PanFS support. For compilers, debugger and application performance analysis tools, the Tri-Laboratory community intends to separately subcontract with the appropriate vendors for support.

#### **4.1 Highly Reliable Management Network (TR-1)**

The SU management Ethernet will be a highly reliable network that does not drop a single node from the network more than once a year. For example an SU design with 162 nodes, the connection between any TLCC2 SU node and the management network will be dropped less than once every ~2,000 months. This is both a hardware and a software (Linux Ethernet device driver) requirement. In addition, the management network will be implemented with connectors on the node mating to the management Ethernet cabling and connectors (Section 3.2.11) so that manually tugging or touching the cable at a node or switch does not drop the Ethernet link. The management Ethernet switches (Section 3.2.11) will be configured such that they behave as standard multi-port bridges. The management Ethernet design will avoid bandwidth oversubscription greater than 16:1 at any point.

Each SU management Ethernet will be connected via one 1Gb/s Ethernet uplink to the RPS node.

#### **4.2 TLCC2 Node Reliability and Monitoring (TR-1)**

The TLCC2 SU nodes and other hardware components will have a Mean Time Between Failure (MTBF) of greater than 217,728 hours (less than one node failure per week per 1,296 nodes). Any failing SU hardware component that causes one or more nodes to becoming unavailable for job scheduling or kills the job running on it is included in this MTBF statistic. For example failures of: DIMMs, processors, motherboards, non-redundant power supplies, blade chassis, IBA infrastructure, PDUs all cause nodes to crash or make nodes become unavailable and are counted in this statistic. Redundant parts that fail such as power supplies, fans, single memory chips that are covered by chip-kill, but do not cause nodes to crash or become unavailable do not count in this statistic. Parts removed under preventive maintenance prior to actual failure such as DIMMs removed after excessive correctable single bit errors or SATA drives removed due to large number of block remaps or SMART future failure indicators do not count in this statistic.

The TLCC2 SU nodes will have real-time hardware monitoring, at a specified interval, of system temperature, processor temperature, fan rotation rate, power supply voltages, etc. This node hardware monitoring facility will alert the scalable monitoring software in Section 4.6 via serial console or management network when any monitored hardware parameter falls outside of the specified nominal range. In addition, the system components may provide failure or diagnostic information via the serial console or management network.

#### **4.3 In Place Node Service (TR-1)**

The SU nodes will be serviceable from within the rack. The node will be mechanically designed to have minimal tool requirements for disassembly. The nodes and other rack components will be mechanically designed so that complete node other rack component disassembly and reassembly can be accomplished in less than 20 minutes by a trained technician without having to move the rack. Blade solutions will have hot swappable blades: the blade chassis will not require being powered down during a blade replacement repair action.

#### **4.4 Component Labeling (TR-1)**

Every rack, Ethernet switch, Ethernet cable, IBA switch, IBA cable, node, disk enclosure will be clearly labeled with a unique identifier visible from the front of the rack and/or the rear of the rack, as appropriate, when the rack door is open. These labels will be high quality so that they do not fall off, fade, disintegrate, or otherwise become unusable or unreadable during the lifetime of the cluster. The font will be non-serif such as Arial or Courier with font size for these labels at least 9pt. Nodes will be labeled from the rear with a unique serial number for inventory tracking. It is desirable that motherboards also have a unique serial number for inventory tracking. This serial number needs to be visible without having to disassemble the node, or else it will be queryable, either from Linux or the BIOS from a Linux command line tool.

#### **4.5 Field Replaceable Unit (FRU) Diagnostics (TR-2)**

Diagnostics will be provided that isolate a failure of a TLCC2 SU component to a single FRU for the supplied hardware. These diagnostics will run from the Linux command line. The diagnostic information will be accessible to operators through networked workstations.

##### **4.5.1 Node Diagnostics Suite (TR-1)**

The Offeror will provide a set of hardware diagnostic programs (a diagnostic suite or diagnostics) for each type of node provided that run from the Linux command line and produce output to STDERR or STDOUT and exit with an appropriate error code when errors are detected. These diagnostics will be capable of stressing the node motherboard components such as processors, chip set, memory hierarchy, on-board networking (e.g., management network), peripheral buses, and local disk drives. The diagnostics will stress the memory hierarchy to generate single and double bit memory errors. The diagnostics will read the hardware single and double bit memory error counters and reset the counts to zero. The diagnostics will stress the local disk in a nondestructive test to generate correctable and uncorrectable read and/or write errors. The diagnostics will read the hardware and/or Linux recoverable I/O error counters and reset the counts to zero. The diagnostics will stress the integer and floating point units in specific core(s) in serial (i.e., one processor and/or HyperThread, as appropriate, at a time) or in parallel (i.e., multiple processors and/or multiple HyperThreads, as appropriate). The CPU stress tests will bind to a specific core and/or HyperThread, as appropriate, by command line option, if possible.

##### **4.5.2 Memory Diagnostics (TR-1)**

The Linux OS will interface to the SU node hardware memory error facility specified in Section 2.3.2 to log all correctable and uncorrectable memory errors on each memory FRU. When the node experiences an uncorrectable memory error, the Linux kernel will report the failing memory FRU and panic the node. The Offeror will provide a Linux utility that can scan the nodes and report correctable and uncorrectable memory errors at the FRU level indicating the exact DIMM location on the motherboard where the failing or failed DIMM is located and reset the counters.

##### **4.5.3 IBA Diagnostics (TR-1)**

The Offeror will provide a set of IBA hardware diagnostic programs (a diagnostic suite or diagnostics) for IBA components provided that run from the Linux command line and produce output to STDERR or STDOUT and exit with an appropriate error code when

errors are detected. These diagnostics will be able to diagnose failures with IBA HCA, cables, and switch hardware down to individual FRU. These diagnostics will run and correctly diagnose failed and intermittently failing hardware within four hours and find **all** failed or intermittently failing components. In addition, these IBA diagnostics will be able to detect slow portions or portions with high bad packet rates of the interconnect infrastructure. IBA diagnostics mentioned above shall be accessible by open source tools such as those provided by the “InfiniBand-diags” open source package.

#### **4.5.4 IPMI Based diagnostics (TR-1)**

If the Offeror proposes IPMI to perform node diagnostics, such as through IPMI sensors, IPMI field replaceable unit (FRU) records, and IPMI system event log (SEL) entries, node diagnostics shall work with FreeIPMI's ipmi-sensors, ipmi-fru, and ipmi-sel respectively. Offeror will publicly release documentation on any OEM motherboard sensors, OEM FRU records, and OEM SEL event data so that they can be interpreted correctly and added for public release in FreeIPMI.

#### **4.5.5 Peripheral Component Diagnostics (TR-2)**

The Offeror will provide a set of hardware diagnostic programs (a diagnostic suite or diagnostics) for each type of peripheral component provided that run from the Linux command line and produce output to STDERR or STDOUT and exit with an appropriate error code when errors are detected. At a minimum, the diagnostics will test the 1000BaseT and other provided networking (e.g., Fibre Channel) adapters, RAID device and disks.

### **4.6 Scalable System Monitoring (TR-2)**

There will be a scalable system monitoring capability for the TLCC2 SU supplied by the Tri-Laboratory Community that has a command line interface (CLI) for scriptable control and monitoring. This facility will directly interface to the Offeror provided node-monitoring facility through the “Motherboard Sensors” Section 3.3.11 facilities. The Tri-Laboratory Community will centrally collect all SU monitoring information at intervals set by the system administrator on one of the service nodes. All SU monitoring and diagnostics information provided by the Offeror will be formatted so that “expect scripts” can detect failures. In addition, this facility will be used by the Tri-Laboratory Community to launch diagnostics in parallel over the management network across all or part of the cluster, as directed by a system administrator from the Linux command line on one of the service nodes.

### **4.7 Hardware Maintenance (TR-1)**

The Offeror will supply hardware maintenance for each proposed TLCC2 SU for a three-year period, which begins with cluster acceptance. Note that this implies that the number of SU under maintenance will ramp up over the delivery schedule and ramp down starting three years after the first TLCC2 cluster is accepted. Tri-Laboratory personnel will attempt on-site first-level hardware fault diagnosis and repair actions. Offeror will provide second-level hardware fault diagnosis and fault determination during normal business hours. That is, if Tri-Laboratory personnel cannot repair failing components with replacements from the on-site parts cache, then the Offeror personnel will be required to make on-site repairs. Offeror supplied hardware maintenance response time will be before the end of the next business day from incident report until Offeror personnel perform diagnosis and/or repair work. The proposed system will be installed in a limited access area vault type rooms (VTR) at the Tri-

Laboratory sites and maintenance personnel must obtain DOE P clearances for repair actions at LLNL and be escorted during repair actions. US Citizenship for maintenance personnel is highly preferred because it takes at least 30 days to obtain VTR access for foreign nationals from non-sensitive countries. During the period from the start of SU installation through acceptance, Offeror support for hardware will be 12 hour a day, seven days a week (0800-2000 PDT for LLNL and 0800-2000 MDT for LANL and Sandia), with one-hour response time.

#### **4.7.1 On-site Parts Cache (TR-1)**

A scalable parts cache (of FRUs and hot spare nodes of each type proposed) at each Tri-Laboratory site is required that will be sufficient to sustain necessary repair actions on all proposed hardware and keep them in fully operational status for at least one week without parts cache refresh. That is, the parts cache, based on Offeror's MTBF estimates for each FRU and each SU, will be sufficient to perform all required repair actions for one week without the need for parts replacement and should be scaled up as SUs are delivered. The Offeror will propose sufficient quantities of FRUs and hot-spare nodes for the parts cache. The parts cache will be enlarged, at the Offeror's expense, should the on-site parts cache prove in actual experience to be insufficient to sustain the actually observed FRU or node failure rates. However, at a minimum, the on-site parts cache will include the following fully configured (except for IBA HCA) nodes: ten (10) compute nodes and two (2) each of LSM, GW and RPS nodes. In addition, a minimum of the following parts (and quantity), if bid: SATA Disks (2), SDRAM DIMM kit for a node (5), power supplies of each type (10), fans of each type (10), management Ethernet switch (1) and TRMS FRUs (1). The Tri-Laboratory Community will administer the nodes as a separate HSC in the unclassified environment. The Tri-Laboratory Community will store and inventory the HSC and other on-site parts cache components. Parts in the parts cache are Government property. Failed parts become Offeror's property when RMAed back to Offeror.

The Offeror will replenish spare parts cache at each Tri-Laboratory site, as parts are consumed, to restore spare parts cache to a level sufficient to sustain necessary repair actions on all proposed hardware and keep them in fully operational status for at least one week. The Offeror will cross-ship replacement parts. That is, the Offeror will ship the replacement part requested prior to receiving the errant part from the Tri-Laboratory site. The Offeror's obligation to replenish spare parts cache at each Tri-Laboratory site will expire three years after the date of final cluster acceptance at the particular Tri-Laboratory site.

FRU with non-volatile memory components (e.g., SATA Disks, remote management cards with flash memory, etc.) where classified data may be stored, cannot be returned to the Offeror via RMA process. Rather, the Tri-Laboratory site must destroy such equipment quickly after removal from the cluster according to DOE security policies and procedures. The Tri-Laboratory site can provide the Offeror with the serial number of the failed FRU or other data about the FRU that Offeror might require, but the actual FRU itself cannot be returned.

#### **4.7.1 Hot Spare Cluster (TR-1)**

The Offeror will provide a Hot Spare Cluster (HSC) that contains the on-site parts cache from section 4.7.1. and provides functions such as hardware burn-in, problem diagnosis,

etc. The Offeror will supply sufficient racks and associated hardware/software to make the HSC a cluster that can run both online and offline diagnostics on every HSC node and the associated components over the management Ethernet. The Offeror will provide software diagnostics that identify failed components and verify functionality of the various system components. This includes CPUs, DIMMs, and Network Cards (both 10Gb and IB). Sufficient IB network infrastructure will be included with the HSC to support connectivity for up to 10 hot-spare nodes. One (1) 1U rack mounted KVM with a slideout display, keyboard, & mouse should be connected to the RPS node in the HSC.

#### **4.7.2 Statement of Volatility (MR)**

The Offeror shall identify any component in the proposed system that persistently holds data in non-volatile memory or storage. Prior to the system and on-site parts cache delivery the Offeror shall provide a Statement of Volatility for all unique FRUs that state whether the FRU contains only volatile memory or storage and thus cannot hold user data after being powered off, or instead it contains non-volatile memory or storage. If it contains the later, the Offeror shall elaborate on the types of data stored and the method by which the data is modified.

### **4.8 Software Support (TR-1)**

The Offeror will supply software maintenance for each Offeror supplied software component, specifically including the supplied BIOS, starting with the first SU acceptance and ending three years after cluster acceptance. Offeror provided software maintenance will include an electronic trouble reporting and tracking mechanism and periodic software updates. In addition, the Offeror will provide software fixes to reported bugs. The electronic trouble reporting and tracking mechanism will allow the Tri-Laboratory community to report bugs and status bug reports 24 hours a day, seven days a week. The Tri-Laboratory community will prioritize software defects so that the Offeror can apply the software maintenance resources to the most important problems. During the period from the start of SU installation through acceptance, Offeror support for supplied software will be 12 hour a day, seven days a week (0800-2000 PDT for LLNL and 0800-2000 MDT for LANL and Sandia), with one hour response time.

### **4.9 Mean Time Between Failure (MTBF) Calculation (TR-1)**

The Offeror will provide the MTBF calculation for each FRU and node type. The Offeror will use these statistics to calculate the MTBF for the provided aggregate TLCC2 SU hardware. This calculation will be performed using a recognized standard. Examples of such standards are Military Standard (Mil Std) 756, Reliability Modeling and Prediction, which can be found in Military Handbook 217F, and the Sum of Parts Method outlined in Bellcore Technical Reference Manual 332. In the absence of relevant technical information in the proposal, the Tri-Laboratory community will be forced to make pessimistic reliability, availability, and serviceability assumptions in evaluating the proposal.

## **End of Section 4**

## 5 Facilities Information

A portion of existing facilities at the Tri-Laboratory sites will be used for siting the TLCC2 SU. The following sections give detailed facilities information about the Tri-Laboratory sites and facilities based requirements for SU and cluster aggregations of SU.

### 5.1 LLNL Facilities Information

At LLNL, SU will be sited in on both the east and west main computer floor in B453 (see Figure 1). This entire facility has approximately 18,000 ft<sup>2</sup> and 2.5 MW of power and associated cooling available for this purpose for the SU compute, networking and peripherals. The computer floor is 48" raised floor with 250 lbs/ft<sup>2</sup> loading capability. However, racks with up to 500-lbs/ft<sup>2</sup> floor loading can be accommodated with additional floor bracing. The overall SU aggregations average floor loading (including isles between rows) cannot exceed 250 lbs/ft<sup>2</sup>. In addition, rolling weight of racks during installation cannot exceed 250 lbs/ft<sup>2</sup>. Power will be provided to racks by under floor electrical outlets supplied by LLNS to Offeror's specifications. Circuit breakers and PDUs are available in wall panels that can be modified to Offeror's specifications. All other cables must be contained in cable trays supplied by LLNS to Offeror's specifications. Straight point-to-point cable runs can **NOT** be assumed. LLNS will provide floor tile cut to Offeror's specifications. In addition, it is anticipated that the Offeror's equipment will be placed in adjacent rows so that air intakes in racks from adjacent rows are abutting with Offeror's specified separations and hot air exhausts in racks from adjacent rack rows are abutting with Offeror's specified separations. That is, the racks will be placed so that there are HOT and COLD aisle ways between racks with chilled air entering in the COLD isles and warmed air exiting in the HOT aisles. Offerors will describe any unique cooling solutions that allow for more efficient utilization of computer floor space and provide information on facilities impacts. LLNS will provide 2'x2' grated floor tiles with 80% void in "cold aisles" to product up to 2,500 CFM airflow per tile.

During installation, racks will transit from Offeror's delivery trucks through 12' (W) x 12' (H) roller doors to an interior delivery dock. Note that the delivery dock height is 45" and can only accommodate one tractor-semitrailer rig at a time. The racks will transit down a 300' long 9' 8" (W) x 8' 6" (H) hallway. The racks must transit several doors of size 7' 10" (W) x 7' 10" (H) and ride a freight elevator up one floor. The freight elevator doors are 8' 4" (W) x 8' (H), the elevator area is 8' (W) x 12' (D) and the maximum loading of the freight elevator is 10,000 lbs. Racks may be staged on the B453 computer floor for unloading from packaging or unpackaged on the interior delivery dock.

This facility will need modifications for siting TLCC2 SU. SU siting costs are considered part of the TCO. LLNS anticipates the construction work to be completed within sixty days of final facilities configuration agreement with Offeror.

### 5.2 LANL Facilities Information

At LANL, most SUs will be sited in the SCC, Building 2327. This entire facility has approximately 20,000 ft<sup>2</sup> and 3.2 MW of power and associated cooling available for this purpose for the SU compute, networking and peripherals. The computer floor is 42" raised floor with 300 lbs/ft<sup>2</sup> loading capability. There is a 16 ft. ceiling and an 18'6" ceiling plenum. The overall SU aggregations average floor loading (including aisles between rows)

cannot exceed 250 lbs/ft<sup>2</sup>. In addition, rolling weight of racks during installation cannot exceed 250 lbs/ft<sup>2</sup>. Power will be provided to racks by under floor electrical outlets supplied by LANL to Offeror's specifications. All other cables must be contained in overhead cable trays to hide them from view. Straight point-to-point cable runs can **NOT** be assumed. LANL will provide floor tile cut to Offeror's specifications. In addition, it is anticipated that the Offeror's equipment will be placed in adjacent rows so that air intakes in racks from adjacent rows are abutting with Offeror's specified separations and hot air exhausts in racks from adjacent rack rows are abutting with Offeror's specified separations. That is, the racks will be placed so that there are HOT and COLD aisle ways between racks with chilled air entering in the 6ft or wider COLD aisles and warmed air exiting in the 4ft wide HOT aisles. The average airflow per floor tile in the SCC is 800-1500 CFM and the cooling limit is 24 kW per rack.

A few SUs delivered to LANL will be sited in the LDCC, Building 1498. This facility has 24" raised computer floor with 250 lb/ft<sup>2</sup> loading capability. There is a 10 ft. ceiling. The average airflow per floor tile in the LDCC is 600-800 CFM. The air-cooled limit is only 10 kW per rack, so higher density liquid cooled solutions are preferred, particularly if the racks are otherwise the same as those proposed for LANL's SCC facility.

Electrical power distribution in both buildings normally uses 208V/120V 3-phase PDUs. Higher voltage such as 480V is possible but only as raw power (not on UPS), so it is not preferred as a total solution.

### 5.3 SNL Facilities Information

At SNL, SUs can be sited in any of the following facilities in New Mexico or California:

SNL New Mexico:

The 880/220 facility has 300 ft<sup>2</sup> available, 400 kW of power, and 130 tons of cooling. There is a 10 ft ceiling and 5 ft ceiling plenum. The overall SU aggregations average floor loading (including aisles between rows) cannot exceed 250lb/ft<sup>2</sup>. Power will be provided by under floor electrical outlets supplied by the Offeror's specifications. Sandia will provide floor tiles cut to the Offeror's specifications. Sandia will provide 2'x2' grated floor tiles with 55% void for the "cold aisles".

The 880/230 facility has 500 ft<sup>2</sup> available, 800 kW of power, and 500 tons of cooling. There is a 10 ft ceiling and 5 ft ceiling plenum. The overall SU aggregations average floor loading (including aisles between rows) cannot exceed 250lb/ft<sup>2</sup>. Power will be provided by under floor electrical outlets supplied by the Offeror's specifications. Sandia will provide floor tiles cut to the Offeror's specifications. Sandia will provide 2'x2' grated floor tiles with 55% void for the "cold aisles".

Electrical power distribution uses 208V/120V single phase or 3-phase PDUs.

Electrical power distribution can also utilize 480V/ 3-phase without PDUs.

SNL California:

At SNL California SU aggregations will be sited in building/room 912/097. The 912/097 facility has 1,750 ft<sup>2</sup>, 800 kW of power and 250 tons of cooling. The computer floor is 18-24" raised floor with 250lb/ft<sup>2</sup> loading capacity. There is a 10 ft ceiling and 2-4 ft ceiling plenum.

The overall SU aggregations average floor loading (including aisles between rows) cannot exceed 250lb/ft<sup>2</sup>. Power will be provided by overhead electrical bus systems supplied by SNL with outlets supplied to the Offeror's specifications. Sandia will provide direct, above the floor, cold air.

#### **5.4 Power Requirements (TR-1)**

Power requirements will be fully disclosed by Offeror at the time of proposal submission. This information will be communicated in the Offeror's proposal. Offeror will provide for each rack type: the number of kW or kVA required, the number and type of power connections required, and anticipated electrical load. This information will be verified by joint written (e-mail and text files) and telephone conferences between the Offeror and the Tri-Laboratory community no more than two months prior to the first SU delivery.

#### **5.5 Cooling Requirements (TR-1)**

Cooling requirements will be fully disclosed by Offeror in the proposal. Offeror will provide for each rack type (the number of BTU or tons AC required) and any environmental requirements, such as temperature and/or humidity range requirements. This information will be verified by joint written (e-mail and text files) and telephone conferences between the Offeror and the Tri-Laboratory community no more than two months prior to the first SU delivery.

#### **5.6 Floor Space Requirements (TR-1)**

Floor space requirements for a cluster with an aggregation of four SU (with associated IBA spine switches) will be fully disclosed by Offeror at the time of proposal submission. This information will be communicated in the Offeror's proposal. Offeror will provide a detailed floor plan (system layout) diagram indicating rack placement and location of required electrical outlets. This information will be verified by joint written (e-mail and text files) and telephone conferences between the Offeror and the Tri-Laboratory community no more than two months prior to the first SU delivery.

#### **5.7 Delivery Requirements (TR-1)**

If Offeror has any delivery requirements these will be communicated to the Tri-Laboratory community in the proposal. TLCC2 SU will be physically located inside a Limited Access Area in a Vault Type Room (VTR). The Tri-Laboratory community will only provide access to the room to authorized personnel under Authorized Escort. All on-site personnel will be required to submit applications for access and be approved by standard Tri-Laboratory procedures prior to entry into this facility. All on-site personnel at LLNL will require being DOE P-cleared or P-clearable. RFP responses should indicate if the on-site team has members that are other than US Citizens and provide the country of citizenship for those members. Physical access to this facility by foreign nationals from sensitive countries will not be allowed. Dialup capability and Internet access to the system will be allowed before the system goes classified, but not afterwards. Authorized individuals may be allowed

## TLCC2 Statement of Work

remote access for running diagnostics and problem resolution only while the system remains unclassified.

Unless otherwise indicated in Offeror's proposal, installation crews will work up to an eight (8) hour day Monday through Friday, 8:00 a.m. to 5:00 p.m. Longer days, differing shift start/end times and/or weekend shifts can be accommodated by the Tri-Laboratory community at Offeror's request at least one week prior to delivery.

### **5.7.1 SU Installation Time (TR-1)**

Offeror will deliver, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over each SU to the Tri-Laboratory community for acceptance testing within **three days** from the time the first truck delivering the SU backs up to the loading dock.

**End of Section 5**

## **6 Project Management**

This procurement envisions purchasing a large number of SUs during 4QCY10 through 3QCY12. The exact number of SUs delivered depends on the peak of the SU. In order to provide DOE/NNSA HQ the maximum amount of budget flexibility, additional SU are required beyond the base SU. The SU deliveries designated with “Option” in them are options that may be exercised at the sole discretion of LLNS upon request by the DOE/NNSA ASC HQ program office or individual laboratory. With delivery of SUs to specific sites, additional 324 port, or larger, spine switches and cables are required as part of this RFP in order to aggregate SU into clusters. The receiving sites are responsible for providing 1-10 Gb/s Ethernet switching infrastructure to integrate the SU into the sites 1-10 Gb/s Ethernet backbones. The receiving sites are also responsible for integrating the delivered SU’s into the sites existing multi-cluster file systems.

The construction, ship testing, delivery, installation, and acceptance testing of the TLCC2 SUs is a complex endeavor. It is anticipated that this project will require close coordination of Tri-Laboratory community, IBA supplier and the Offeror’s personnel.

### **6.1 Risk Reduction Plan (TR-1)**

The Offeror will provide a list of the top risks associated with this project from their point of view. The list should be categorized by their impact (low, medium, high) and probability of occurrence (low, medium, high). For each risk, the Offeror will state the plan to mitigate and/or the alternative solution capable of delivering success on schedule.

### **6.2 Open Source Development Partnership (TR-2)**

The Offeror will provide information on the capabilities of the Offeror to engage in an Open Source development partnership and meet the goals set out in Section 1, 2, and 3 (i.e., OpenFabrics, Free IPMI, OpenMPI). This information should include Offeror’s financial health; Offeror’s qualifications as a cluster provider; Offeror’s qualifications as an Open Source development organization; cluster product roadmap and comparison to the overall TLCC2 strategy; the willingness of the Offeror to participate in the Open Source development, with other partners, of key missing HPTC cluster technology components such as scalable parallel file systems and cluster resource scheduling. If the Offeror has technology, such as a scalable parallel file system, cluster management tools, or cluster resource scheduling, that could be contributed to the Open Source community, please indicate that as well in the proposal.

### **6.3 Project Manager (TR-1)**

The Offeror will provide the name and resume of the proposed project manager within the Offeror’s corporation for the proposed activity. This project manager will be approved by LLNS technical representative. The project manager must be empowered by the Offeror’s corporation to plan and execute the construction, shipment, and installation of the proposed configuration. This must include sufficient personnel and hardware resources within the corporation to assure successful completion of the activity on the proposed schedule. The project manager must be empowered by the Offeror's corporation to facilitate and/or coordinate timely BIOS, firmware, and software fixes or updates. This must include sufficient access to engineering personnel and expertise to assure successful completion of

the activity on the proposed schedule. Offeror must be empowered to facilitate and/or coordinate with vendor partner corporations as well as their own.

## 6.4 Project Milestones (TR-3)

Delivery of the N SUs may be accomplished within three quarter year periods starting no later than 3QCY2011 and ending no later than 1QCY2012. Recall that N is defined in section 3.2.4. Let  $X = \text{int}(N/3)$  and  $Y = N - 2 * X$ .

Example: for  $N = 43$  (50.0 teraFLOP/s SU), then  $X = 14$  and  $Y = 15$ .

For the purposes of this section, the Tri-Laboratory community assumes a delivery schedule as follows:

Delivery Phase	Quantity of SUs	Delivery Quarter
1	Y SU	3QCY2011
2	X SU	4QCY2011
3	X SU	1QCY2012
Total	N SU	

The Offeror may propose alternative SU delivery schedules within the start of the subcontract to the end of 1QCY2012.

Offeror will deliver, additional spine switches, cables to enable the Tri-Laboratory personnel to combine multiple SU into larger clusters, as directed by the Tri-Laboratory in writing 30 days prior to SU delivery.

The Tri-Laboratory community plans to transition the TLCC2 SUs aggregated into clusters to classified operation Limited Availability (LA) Production status for a small group of ASC Program and SSP with fifteen days of delivery. The Tri-Laboratory community plans to run in LA status for a week and then transition the TLCC2 cluster to General Availability (GA) for the general community of ASC Program and SSP users. Instability of SU during acceptance, transition to LA or transition to GA, will impede this schedule and should be avoided: time is of the essence.

In addition, there are multiple activities among multiple institutions and organizations within the institutions that must be coordinated prior to the first delivery and ongoing during the seven quarters of SU deliveries. In order to assure the timely execution of these programmatic goals and to make sure both parties understand the timeline, the Offeror will provide the Tri-Laboratory community with a project plan no more than seven days after subcontract award.

### 6.4.1 Detailed Project Plan (TR-1)

The Offeror will provide a detailed project plan no more than seven days after subcontract award. This project plan will include a Gantt chart with all the project milestones with dates and durations for work activities leading up to the milestones. The Gantt chart will indicate work activity and milestone and organizational dependencies. The Gantt chart will clearly indicate the project's critical path. At least one level of detail below each of the project milestones showing the work activities leading up to completion of the milestone will be included in the Gantt chart. The project plan will include a written Tri-Laboratory

TOSS build image checkout plan, pre-ship test plan and acceptance test plan. The project plan Gantt chart will be a Microsoft Project data file. The test plans will be Microsoft Word data files.

The build image checkout, pre-ship test and acceptance test plans will be mutually agreeable, but will include:

- Successfully running with correct results three mixed MPI/OpenMP jobs (sPPM, UMT2K, LINPACK) sequentially or simultaneously across 90% of the SU compute nodes for at least four hours without failure;
- Successfully running the LLNL Presta MPI stress test sequentially or simultaneously across 90% of the SU compute nodes for four hours without failure or performance anomalies; and
- A demonstration that the Management Ethernet is functional, stable, and reliable.

Offeror will be responsible for LINPACK tuning and execution. The Tri-Laboratory community will be responsible for sPPM and UMT2K tuning and execution. The test plans will include clear test entry and exit criteria as well as a list of testing activities and benchmarks.

As a part of this detailed project plan, the Offeror will provide an updated and detailed Risk Reduction Plan (initially presented in Section 6.1). This plan will be developed jointly, in partnership with the Tri-Lab community, and evolve in response to issues identified throughout the life of the project.

This milestone is complete when the LLNS Contract Administrator (LCA) approves the project plan.

#### **6.4.2 Tri-Laboratory TOSS Final Checkout (TR-1, April 15, 2011)**

The Tri-Laboratory community will provide TOSS (section 2.3.1) for installation on SU. The Offeror will assist the Tri-Laboratory community in development of provided components, test, debug and installation of these software stacks. This TOSS Linux Build Image effort will commence upon subcontract signing and continue throughout the subcontract time span. Offeror will assist the Tri-Laboratory community to finalize these software stacks prior to each SU manufacture. TOSS will then be used to manufacture, test, deliver and accept the SU. This milestone is complete when the first Offeror and Tri-Laboratory joint testing of the TOSS Build Image completes the checkout test plan exit criteria by April 15, 2011.

#### **6.4.3 TLCC2 Phase 1 Build (TR-1, June 2011)**

The Offeror will build, fully assemble, configure, burn-in and test the Y SU for TLCC2 Phase 1 defined in section 6.4, as bid, with Tri-Laboratory Linux Build Images as directed by the LCA.

Offeror shall burn in and stress test TLCC2 Phase 1 equipment (including IBA 4x QDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x QDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software

modifications made by Offeror to successfully complete this 48-hour stress test shall be approved by the LCA. Offeror shall demonstrate that the IBA 4x QDR interconnect with this TLCC2 Phase 1 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC2 Phase 1 SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x QDR L1 infrastructure must be functional); 2) LLNS confirms that the correct TOSS Build Image is installed on TLCC2 Phase 1 SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) LLNS confirms that TLCC2 Phase 1 SU meets the SWL pre-ship test entry criteria; 5) SWL pre-ship test is successfully executed on the TLCC2 Phase 1 SU; 6) the TLCC2 Phase 1 SU successfully completes the SWL pre-ship test exit criteria; 7) the LCA authorizes shipment of SU to tri-Laboratory sites; 8) all equipment for this milestone leaves Offeror's integration location; and 9) the required documentation is approved by the LCA.

### **6.4.4 TLCC2 SU Phase 1 Delivery and Acceptance (TR-1, July 2011)**

Offeror will deliver the TLCC2 SU to the Tri-Laboratory sites as directed by the LCA, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over the TLCC2 SU to the Tri-Laboratory community for acceptance testing. In addition, Offeror will deliver sufficient IBA spine switches and cables, as directed by the LCA, in order to allow the Tri-Laboratory community to assemble multiple SU clusters.

Offeror will burn in and stress test TLCC2 Phase 1 equipment (including IBA 4x QDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x QDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48-hour stress test shall be approved by the LCA. Offeror shall demonstrate that the IBA 4x QDR interconnect with this TLCC2 Phase 1 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC2 Phase 1 SU with IBA hardware is installed at the Tri-Laboratory sites, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x QDR L1 infrastructure must be functional); 2) the LCA confirms that the correct TOSS Build Image is installed on the TLCC2 Phase 1 SUs; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) LLNS confirms that TLCC2 Phase 1 SUs meet the SWL post-ship test entry criteria; 5) SWL post-ship test is successfully executed on the TLCC2 Phase 1 SUs; 6) the TLCC2 Phase 1 SUs successfully completes the SWL post-ship test exit criteria; and 7) the required documentation is approved by the LCA.

### **6.4.5 TLCC2 Phase 1 Cluster Integration (TR-1, July 2011)**

Offeror will deliver Phase 1 on-site hardware maintenance parts cache to each Phase 1 Tri-Laboratory site. Offeror will integrate the site specific SU's into the single multi-SU (2, 4,

or 8 SU) fully functional Clusters. The Offeror will fully assemble, pass Tri-Laboratories' delivery checklist and initial functionality and performance verification testing, and turn over the integrated multi-SU Clusters to the receiving Laboratory for integration testing.

Completion of this milestone starts the three-year maintenance clock on the TLCC2 Phase 1 clusters.

This milestone is complete when: 1) all Cluster SU and IBA spine switch and cabling hardware is installed, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, the L1 and L2 IBA infrastructure must be functional); 2) LLNS confirms that the correct software stack is installed on the Cluster; 3) the LCA confirms that the Cluster on-site hardware maintenance parts cache is fully stocked; 4) the LCA confirms that the Cluster meets the integration test entry criteria; 5) the acceptance test is successfully executed on the Cluster; 6) the Cluster successfully completes the integration test exit criteria; and 7) the required documentation is approved by the LCA.

#### **6.4.6 TLCC2 Phase 1 Option Build (TR-1, July 2011)**

At the Option of LLNS, the Offeror will build, fully assemble, configure, burn-in and test up to an additional 4 SU for TLCC2 Phase 1 Option, as bid, with Tri-Laboratory Linux Build Images as directed by the LCA.

Offeror shall burn in and stress test TLCC2 Phase 1 Option equipment (including IBA 4x QDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x QDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48-hour stress test shall be approved by the LCA. Offeror shall demonstrate that the IBA 4x QDR interconnect with this TLCC2 Phase 1 Option cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC2 Phase 1 Option SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x QDR L1 infrastructure must be functional); 2) LLNS confirms that the correct TOSS Build Image is installed on TLCC2 Phase 1 Option SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) LLNS confirms that TLCC2 Phase 1 Option SU meets the SWL pre-ship test entry criteria; 5) SWL pre-ship test is successfully executed on the TLCC2 Phase 1 Option SU; 6) The TLCC2 Phase 1 Option SU successfully completes the SWL pre-ship test exit criteria; 7) the LCA authorizes shipment of SU to Tri-Laboratory sites; 8) all equipment for this milestone leaves Offeror's integration location; and 9) the required documentation is approved by the LCA.

#### **6.4.7 TLCC2 SU Phase 1 Option Delivery and Acceptance (TR-1, August 2011)**

Offeror will deliver the TLCC2 SU to the Tri-Laboratory sites as directed by the LCA, install, fully assemble, pass Offeror's delivery check list and initial functionality and

performance verification testing, and turn over the TLCC2 SU to the Tri-Laboratory community for acceptance testing. In addition, Offeror will deliver sufficient IBA spine switches and cables, as directed by the LCA, in order to allow the Tri-Laboratory community to assemble multiple SU clusters.

Offeror will burn in and stress test TLCC2 Phase 1 Option equipment (including IBA 4x QDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x QDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48-hour stress test shall be approved by the LCA. Offeror shall demonstrate that the IBA 4x QDR interconnect with this TLCC2 Phase 1 Option cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC2 Phase 1 Option SU with IBA hardware is installed at the tri-Laboratory sites, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x QDR L1 infrastructure must be functional); 2) the LCA confirms that the correct TOSS Build Image is installed on the TLCC2 Phase 1 Option SUs; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) LLNS confirms that TLCC2 Phase 1 Option SUs meet the SWL post-ship test entry criteria; 5) SWL post-ship test is successfully executed on the TLCC2 Phase 1 Option SUs; 6) the TLCC2 Phase 1 Option SUs successfully completes the SWL post-ship test exit criteria; and 7) the required documentation is approved by the LCA.

### **6.4.8 TLCC2 Phase 1 Option Cluster Integration (TR-1, August 2011)**

Offeror will deliver Phase 1 Option on-site hardware maintenance parts cache to each Phase 1 Option Tri-Laboratory site. Offeror will integrate the site specific SU's into the single multi-SU (2, 4, or 8 SU) fully functional Clusters. The Offeror will fully assemble, pass Tri-Laboratories' delivery checklist and initial functionality and performance verification testing, and turn over the integrated multi-SU Clusters to the receiving Laboratory for integration testing.

Completion of this milestone starts the three-year maintenance clock on the TLCC2 Phase 1 Option clusters.

This milestone is complete when: 1) all Cluster SU and IBA spine switch and cabling hardware is installed, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, the L1 and L2 IBA infrastructure must be functional); 2) LLNS confirms that the correct Laboratory software stack (see Section 2.3) is installed on the Cluster; 3) the LCA confirms that the Cluster on-site hardware maintenance parts cache is fully stocked; 4) the LCA confirms that the Cluster meets the integration test entry criteria; 5) the acceptance test is successfully executed on the Cluster; 6) the Cluster successfully completes the integration test exit criteria; and 8) the required documentation is approved by the LCA.

#### **6.4.9 TLCC2 Phase 2 SU Build (TR-1, August 2011)**

The Offeror will build, fully assemble, configure, burn-in and test the X SU for TLCC2 Phase 2 defined in section 6.4, as bid, with Tri-Laboratory Linux Build Images as directed by the LCA.

Offeror shall burn in and stress test TLCC2 Phase 2 equipment (including IBA 4x QDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x QDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48-hour stress test shall be approved by the LCA. Offeror shall demonstrate that the IBA 4x QDR interconnect with this TLCC2 Phase 2 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC2 Phase 2 SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x QDR L1 infrastructure must be functional); 2) LLNS confirms that the correct TOSS Build Image is installed on TLCC2 Phase 2 SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) LLNS confirms that TLCC2 Phase 2 SU meets the SWL pre-ship test entry criteria; 5) SWL pre-ship test is successfully executed on the TLCC2 Phase 2 SU; 6) the TLCC2 Phase 2 SU successfully completes the SWL pre-ship test exit criteria; 7) the LCA authorizes shipment of SU to tri-Laboratory sites; 8) all equipment for this milestone leaves Offeror's integration location; and 9) the required documentation is approved by the LCA.

#### **6.4.10 TLCC2 Phase 2 SU Delivery and Acceptance (TR-1, September 2011)**

Offeror will deliver the TLCC2 SU to the Tri-Laboratory sites as directed by the LCA, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over the TLCC2 SU to the Tri-Laboratory community for acceptance testing. In addition, Offeror will deliver sufficient IBA spine switches and cables, as directed by the LCA, in order to allow the Tri-Laboratory community to assemble multiple SU clusters.

Offeror will burn in and stress test TLCC2 Phase 2 equipment (including IBA 4x QDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x QDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48-hour stress test shall be approved by the LCA. Offeror shall demonstrate that the IBA 4x QDR interconnect with this TLCC2 Phase 2 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

## TLCC2 Statement of Work

This milestone is complete when: 1) the TLCC2 Phase 2 SU with IBA hardware is installed at the Tri-Laboratory sites, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x QDR L1 infrastructure must be functional); 2) the LCA confirms that the correct TOSS Build Image is installed on the TLCC2 Phase 2 SUs; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) LLNS confirms that TLCC2 Phase 2 SUs meet the SWL post-ship test entry criteria; 5) SWL post-ship test is successfully executed on the TLCC2 Phase 2 SUs; 6) the TLCC2 Phase 2 SUs successfully completes the SWL post-ship test exit criteria; and 7) the required documentation is approved by the LCA.

### **6.4.11 TLCC2 Phase 2 Cluster Integration (TR-1, September 2011)**

Offeror will deliver Phase 2 on-site hardware maintenance parts cache to each Phase 2 tri-Laboratory site. Offeror will integrate the site specific SU's into the single multi-SU (2, 4, or 8 SU) fully functional Clusters. The Offeror will fully assemble, pass Tri-Laboratories' delivery checklist and initial functionality and performance verification testing, and turn over the integrated multi-SU Clusters to the receiving Laboratory for integration testing.

Completion of this milestone starts the three-year maintenance clock on the TLCC2 Phase 2 clusters.

This milestone is complete when: 1) all Cluster SU and IBA spine switch and cabling hardware is installed, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, the L1 and L2 IBA infrastructure must be functional); 2) LLNS confirms that the correct software stack is installed on the Cluster; 3) the LCA confirms that the Cluster on-site hardware maintenance parts cache is fully stocked; 4) the LCA confirms that the Cluster meets the integration test entry criteria; 5) the acceptance test is successfully executed on the Cluster; 6) the Cluster successfully completes the integration test exit criteria; and 7) the required documentation is approved by the LCA.

### **6.4.12 TLCC2 Phase 3 Build (TR-1, October 2011)**

The Offeror will build, fully assemble, configure, burn-in and test the Y SU for TLCC2 Phase 3 defined in section 6.4, as bid, with Tri-Laboratory Linux Build Images as directed by the LCA.

Offeror shall burn in and stress test TLCC2 Phase 3 equipment (including IBA 4x QDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x QDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48-hour stress test shall be approved by the LCA. Offeror shall demonstrate that the IBA 4x QDR interconnect with this TLCC2 Phase 3 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC2 Phase 3 SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x QDR L1 infrastructure must

be functional); 2) LLNS confirms that the correct TOSS Build Image is installed on TLCC2 Phase 3 SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) LLNS confirms that TLCC2 Phase 3 SU meets the SWL pre-ship test entry criteria; 5) SWL pre-ship test is successfully executed on the TLCC2 Phase 3 SU; 6) The TLCC2 Phase 3 SU successfully completes the SWL pre-ship test exit criteria; 7) the LCA authorizes shipment of SU to tri-Laboratory sites; 8) all equipment for this milestone leaves Offeror's integration location; and 9) the required documentation is approved by the LCA.

#### **6.4.13 TLCC2 SU Phase 3 Delivery and Acceptance (TR-1, November 2011)**

Offeror will deliver the TLCC2 SU to the Tri-Laboratory sites as directed by the LCA, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over the TLCC2 SU to the Tri-Laboratory community for acceptance testing. In addition, Offeror will deliver sufficient IBA spine switches and cables, as directed by the LCA, in order to allow the Tri-Laboratory community to assemble multiple SU clusters.

Offeror will burn in and stress test TLCC2 Phase 3 equipment (including IBA 4x QDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x QDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48-hour stress test shall be approved by the LCA. Offeror shall demonstrate that the IBA 4x QDR interconnect with this TLCC2 Phase 3 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC2 Phase 3 SU with IBA hardware is installed at the Tri-Laboratory sites, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x QDR L1 infrastructure must be functional); 2) the LCA confirms that the correct TOSS Build Image is installed on the TLCC2 Phase 3 SUs; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) LLNS confirms that TLCC2 Phase 3 SUs meet the SWL post-ship test entry criteria; 5) SWL post-ship test is successfully executed on the TLCC2 Phase 3 SUs; 6) the TLCC2 Phase 3 SUs successfully completes the SWL post-ship test exit criteria; and 7) the required documentation is approved by the LCA.

#### **6.4.14 TLCC2 Phase 3 Cluster Integration (TR-1, November 2011)**

Offeror will deliver Phase 3 on-site hardware maintenance parts cache to each Phase 3 Tri-Laboratory site. Offeror will integrate the site specific SU's into the single multi-SU (2, 4, or 8 SU) fully functional Clusters. The Offeror will fully assemble, pass Tri-Laboratories' delivery checklist and initial functionality and performance verification testing, and turn over the integrated multi-SU Clusters to the receiving Laboratory for integration testing.

Offeror shall burn in and stress test TLCC2 Phase 3 equipment (including IBA 4x QDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable

## TLCC2 Statement of Work

Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x QDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48-hour stress test shall be approved by the LCA. Offeror shall demonstrate that the IBA 4x QDR interconnect with this TLCC2 Phase 3 cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

Completion of this milestone starts the three-year maintenance clock on the TLCC2 Phase 3 clusters.

This milestone is complete when: 1) the TLCC2 Phase 3 SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x QDR L1 infrastructure must be functional); 2) LLNS confirms that the correct TOSS Build Image is installed on TLCC2 Phase 3 SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) LLNS confirms that TLCC2 Phase 3 SU meets the SWL pre-ship test entry criteria; 5) SWL pre-ship test is successfully executed on the TLCC2 Phase 3 SU; 6) The TLCC2 Phase 3 SU successfully completes the SWL pre-ship test exit criteria; 7) the LCA authorizes shipment of SU to tri-Laboratory sites; 8) all equipment for this milestone leaves Offeror's integration location; and 9) the required documentation is approved by the LCA.

### **6.4.15 TLCC2 Phase 3 Option Build (TR-1, November 2011)**

At the Option of LLNS, the Offeror will build, fully assemble, configure, burn-in and test up to an additional 4 SU for TLCC2 Phase 3 Option, as bid, with Tri-Laboratory Linux Build Images as directed by the LCA.

Offeror shall burn in and stress test TLCC2 Phase 3 Option equipment (including IBA 4x QDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x QDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48-hour stress test shall be approved by the LCA. Offeror shall demonstrate that the IBA 4x QDR interconnect with this TLCC2 Phase 3 Option cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC2 Phase 3 Option SU with IBA hardware is installed at Offeror's integration location, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x QDR L1 infrastructure must be functional); 2) LLNS confirms that the correct TOSS Build Image is installed on TLCC2 Phase 3 Option SU; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) LLNS confirms that TLCC2 Phase 3 Option SU meets the SWL pre-ship test entry criteria; 5) SWL pre-ship test is successfully executed on the TLCC2 Phase 3 Option SU; 6) the TLCC2 Phase 3 Option SU successfully completes the SWL pre-ship test exit criteria; 7) the LCA authorizes shipment of SU to Tri-Laboratory sites; 8) all

equipment for this milestone leaves Offeror's integration location; and 9) the required documentation is approved by the LCA.

#### **6.4.16 TLCC2 SU Phase 3 Option Delivery and Acceptance (TR-1, December 2011)**

Offeror will deliver the TLCC2 SU to the Tri-Laboratory sites as directed by the LCA, install, fully assemble, pass Offeror's delivery check list and initial functionality and performance verification testing, and turn over the TLCC2 SU to the Tri-Laboratory community for acceptance testing. In addition, Offeror will deliver sufficient IBA spine switches and cables, as directed by the LCA, in order to allow the Tri-Laboratory community to assemble multiple SU clusters.

Offeror will burn in and stress test TLCC2 Phase 3 Option equipment (including IBA 4x QDR interconnect) and replace failing hardware and continue burn in and stress testing of the equipment until the early life failure rate is below one node and/or IBA Field Replaceable Unit (FRU) failure per 48 hour period. After passing Offeror burn in and initial stress tests, Offeror shall stress test the SU IBA 4x QDR interconnect for at least 48 hours without hardware fabric errors or uncovering hardware or software bugs. Any software modifications made by Offeror to successfully complete this 48-hour stress test shall be approved by the LCA. Offeror shall demonstrate that the IBA 4x QDR interconnect with this TLCC2 Phase 3 Option cluster is fully functional and error free with the execution of the SWL for at least five (5) days without any IBA hardware errors.

This milestone is complete when: 1) the TLCC2 Phase 3 Option SU with IBA hardware is installed at the Tri-Laboratory sites, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, IBA 4x QDR L1 infrastructure must be functional); 2) the LCA confirms that the correct TOSS Build Image is installed on the TLCC2 Phase 3 Option SUs; 3) the Offeror successfully completes the IBA interconnect 48 hour stress test; 4) LLNS confirms that TLCC2 Phase 3 Option SUs meet the SWL post-ship test entry criteria; 5) SWL post-ship test is successfully executed on the TLCC2 Phase 3 Option SUs; 6) the TLCC2 Phase 3 Option SUs successfully completes the SWL post-ship test exit criteria; and 7) the required documentation is approved by the LCA.

#### **6.4.17 TLCC2 Phase 3 Option Cluster Integration (TR-1, December 2011)**

Offeror will deliver Phase 3 Option on-site hardware maintenance parts cache to each Phase 3 Option Tri-Laboratory site. Offeror will integrate the site specific SU's into the single multi-SU (2, 4, or 8 SU) fully functional Clusters. The Offeror will fully assemble, pass University's delivery check list and initial functionality and performance verification testing, and turn over the integrated multi-SU Clusters to LLNS for integration testing.

Completion of this milestone starts the three-year maintenance clock on the TLCC2 Phase 3 option clusters.

This milestone is complete when: 1) all Cluster SU and IBA spine switch and cabling hardware is installed, burned-in, and functional (all nodes must be functional, management Ethernet must be functional, the L1 and L2 IBA infrastructure must be functional); 2) LLNS confirms that the correct software stack is installed on the Cluster; 3) the LCA

## TLCC2 Statement of Work

confirms that the Cluster on-site hardware maintenance parts cache is fully stocked; 4) the LCA confirms that the Cluster meets the integration test entry criteria; 5) the acceptance test is successfully executed on the Cluster; 6) the Cluster successfully completes the integration test exit criteria; and 7) the required documentation is approved by the LCA.

### **End of Section 6**

## 7 Glossary

### 7.1 General

<b>Mandatory requirements designated as (MR)</b>	Mandatory Requirements (designated MR) in the Draft Statement of Work (SOW) are performance features that are essential to Tri-Laboratory requirements. An Offeror must satisfactorily propose all Mandatory Requirements in order to have its proposal considered responsive.
<b>Target Requirements designated as (TR-1, TR-2 and TR-3)</b>	Target Requirements (designated TR-1, TR-2, or TR-3), identified throughout the Draft SOW, are features, components, performance characteristics, or other properties that are important to the Tri-Laboratory. However, omission of a response for a Target Requirement will not render a proposal non-responsive. Target Requirements add value to a proposal. Target Requirements are prioritized by dash number. TR-1 is most desirable to the Tri-Laboratory, while TR-2 is more desirable than TR-3. Target Requirement responses will be considered as part of the proposal evaluation process.
<b>Mandatory Option Requirement (MOR)</b>	Mandatory Option Requirement (designated MOR) in the Draft SOW reflects a particular Scalable Unit (SU) configuration required by LANL. LANL needs the ability to acquire this SU configuration as an option. An Offeror must satisfactorily propose all MOR in order to have its proposal considered eligible for award of a subcontract for LANL SUs.

### 7.2 Hardware

<b>b</b>	bit. A single, indivisible binary unit of electronic information.
<b>B</b>	Byte. A collection of eight (8) bits.
<b>32b floating-point arithmetic</b>	Executable binaries (user applications) with 32b (4B) floating-point number representation and arithmetic. Note that this is independent of the number of bytes (4 or 8) utilized for memory reference addressing.
<b>32b virtual memory addressing</b>	All virtual memory addresses in a user application are 32b (4B) integers. Note that this is independent of the type of floating-point number representation and arithmetic.
<b>64b floating-point arithmetic</b>	Executable binaries (user applications) with 64b (8B) floating-point number representation and arithmetic. Note that this is independent of the number of bytes (4 or 8) utilized for memory reference addressing.
<b>64b virtual memory addressing</b>	All virtual memory addresses in a user application are 64b (8B) integers. Note that this is independent of the type of floating-point number representation and arithmetic. Note that all user applications should be compiled, loaded with Offeror supplied libraries and executed with 64b virtual memory addressing by default.
<b>CE</b>	On-site hardware customer engineer performing hardware installation or maintenance (with DOE P-clearance for LLNL).

## TLCC2 Statement of Work

<b>Cluster</b>	A set of SMPs connected via a scalable network technology. The network will support high bandwidth, low latency message passing. It will also support remote memory referencing.
<b>CPU or core or processor</b>	Central Processing Unit or “core” or processor. A VLSI chip constituting one or more computational core(s) (integer, floating point, and branch units), registers and memory interface (virtual memory translation, TLB, and bus controller) and associated cache.
<b>FLOP or OP</b>	Floating Point Operation.
<b>FLOPS or OPS</b>	Plural of FLOP.
<b>FLOP/s or OP/s</b>	Floating Point Operation per second.
<b>FRU</b>	Field Replaceable Unit (FRU) is an aggregation of parts that is a single unit and can be replaced upon failure in the field.
<b>FSB</b>	Front-side bus
<b>GB</b>	gigaByte. gigaByte is a billion base 10 bytes. This is typically used in every context except for Random Access Memory size and is $10^9$ (or 1,000,000,000) bytes.
<b>GiB</b>	gibiByte. gibibyte is a billion base 2 bytes. This is typically used in terms of Random Access Memory and is $2^{30}$ (or 1,073,741,824) bytes. For a complete description of SI units for prefixing binary multiples see URL: <a href="http://physics.nist.gov/cuu/Units/binary.html">http://physics.nist.gov/cuu/Units/binary.html</a> .
<b>GFE</b>	Government Furnished Equipment (GFE) is equipment supplied to the Offeror by the Tri-Laboratory’s when TLCC2 SU build or installation takes place.
<b>GFLOP/s or GOP/s</b>	gigaFLOP/s. Billion ( $10^9=1,000,000,000$ ) 64-bit floating-point operations per second.
<b>HSC</b>	Hot Spare Cluster. A set of nodes on-site at LLNL, LANL and SNL that can be used as a hot spare pool constructed as a stand-alone cluster. This HSC will be used to run diagnostics on failing nodes (after they are swapped out of TLCC2) to determine root cause for failures and to potentially test software releases.
<b>IBA</b>	InfiniBand Architecture see <a href="http://www.openfabrics.org">http://www.openfabrics.org</a> and <a href="http://www.infinibandta.org">http://www.infinibandta.org</a>
<b>IPMI</b>	Intelligent Platform Management Interface. See <a href="http://www.intel.com/design/servers/ipmi/">http://www.intel.com/design/servers/ipmi/</a>
<b>ISA</b>	Instruction Set Architecture
<b>MB</b>	megaByte. megaByte is a million base 10 bytes. This is typically used in every context except for Random Access Memory size and is $10^6$ (or 1,000,000) bytes.
<b>MiB</b>	mebiByte. mebibyte is a million base 2 bytes. This is typically used in terms of Random Access Memory and is $2^{20}$ (or 1,048,576) bytes. For a complete description of SI units for prefixing binary multiples see URL: <a href="http://physics.nist.gov/cuu/Units/binary.html">http://physics.nist.gov/cuu/Units/binary.html</a>
<b>MDS</b>	Lustre Meta Data Server. Performs the Lustre file system functions associated with file system layout and name space mapping.
<b>MFLOP/s or MOP/s</b>	megaFLOP/s. Million ( $10^6=1,000,000$ ) 64-bit floating-point operations per second.

## TLCC2 Statement of Work

<b>MTBF</b>	Mean Time Between Failure. A measurement of the expected reliability of the system or component. The MTBF figure can be developed as the result of intensive testing, based on actual product experience, or predicted by analyzing known factors. See URL: <a href="http://www.t-cubed.com/faq_mtbef.htm">http://www.t-cubed.com/faq_mtbef.htm</a>
<b>Node</b>	Four socket AMD x86-64 or Intel EM64T (or binary compatible) quad core die in an SMP configuration with the Linux operating system and IBA HCA.
<b>OSS</b>	Lustre Object Storage Server. The hardware and software associated with the Lustre Object Storage Targets. OSS connects to TLCC2 via 10 Gb/s Ethernet.
<b>PCIe2 x8</b>	The PCIe Gen 2 standard with 8 lanes of electrically live links. It is not acceptable to have a x8 slot with a x4 electrical connection.
<b>PDU</b>	Power Distribution Unit. Mechanism by which power is distributed to nodes from the higher amperage wall panel.
<b>Peak Rate</b>	The maximum number of 64-bit floating-point instructions (add, subtract, multiply or divide) per second that could conceivably be retired by the system. For microprocessors the peak rate is typically calculated as the maximum number of floating point instructions retired per clock times the clock rate.
<b>POST</b>	Power-On Self Test (POST) is a set of diagnostics that run when the node is powered on to detect all hardware components and verify correct functioning.
<b>SPC</b>	Serial Port Concentrator (SPC) is a rack-mounted device (that may be combined with the RPC) that connects the serial ports of nodes to the management Ethernet via reverse telnet protocol. This allows system administrators to log into the serial port of every node via the management network and perform management actions on the node. In addition, this interface allows the system administrators to set up telnet sessions with each node and log all console traffic.
<b>Scalable</b>	A system attribute that increases in performance or size as some function of the peak rating of the system. The scaling regime of interest is at least within the range of 1 teraFLOP/s to 60.0 (and possibly to 120.0) teraFLOP/s peak rate.
<b>SMP</b>	Shared memory Multi-Processor. A set of CPUs sharing random access memory within the same memory address space. The CPUs are connected via a high speed, low latency mechanism to the set of hierarchical memory components. The memory hierarchy consists of at least processor registers, cache and memory. The cache will also be hierarchical. If there are multiple caches, they will be kept coherent automatically by the hardware. The main memory will be UMA architecture. The access mechanism to every memory element will be the same from every processor. More specifically, all memory operations are done with load/store instructions issued by the CPU to move data to/from registers from/to the memory.
<b>SU</b>	Scalable Unit (SU) is the (nearly) identical replicate unit of hardware envisioned by this statement of work.

<b>Tera-Scale</b>	The environment required to fully support production-level, realized teraFLOP/s performance. This environment includes a robust and balanced processor, memory, mass storage, I/O, and communications subsystems; robust code development environment, tools and operating systems; and an integrated cluster wide systems management and full system reliability and availability.
<b>TB</b>	TeraByte. TeraByte is a trillion base 10 bytes. This is typically used in every context except for Random Access Memory size and is $10^{12}$ (or 1,000,000,000,000) bytes.
<b>TiB</b>	TebiByte. TebiByte is a trillion bytes base 2 bytes. This is typically used in terms of Random Access Memory and is $2^{40}$ (or 1,099,511,627,776) bytes. For a complete description of SI units for prefixing binary multiples see URL: <a href="http://physics.nist.gov/cuu/Units/binary.html">http://physics.nist.gov/cuu/Units/binary.html</a>
<b>TFLOP/s</b>	teraFLOP/s. Trillion ( $10^{12}$ =1,000,000,000,000) 64-bit floating-point operations per second.
<b>UMA</b>	Uniform Memory Access architecture. The distance in processor clocks between processor registers and every element of main memory is the same. That is, a load/store operation has the same latency, no matter where the target location is in main memory.

### 7.3 Software

<b>32b executable</b>	Executable binaries (user applications) with 32b (4B) virtual memory addressing. Note that this is independent of the number of bytes (4 or 8) utilized for floating-point number representation and arithmetic.
<b>64b executable</b>	Executable binaries (user applications) with 64b (8B) virtual memory addressing. Note that this is independent of the number of bytes (4 or 8) utilized for floating-point number representation and arithmetic. Note that all user applications should be compiled, loaded with Offeror supplied libraries and executed with 64b virtual memory addressing by default.
<b>API</b> (Application Programming Interface)	Syntax and semantics for invoking services from within an executing application. All APIs will be available to both Fortran and C programs, although implementation issues (such as whether the Fortran routines are simply wrappers for calling C routines) are up to the supplier.
<b>BIOS</b>	Basic Input-Output System (BIOS) is low level (typically assembly language) code usually held in flash memory on the node that tests and functions the hardware upon power-up or reset or reboot and loads the operating system.
<b>Current standard</b>	Term applied when an API is not “frozen” on a particular version of a standard, but will be upgraded automatically by Offeror as new specifications are released (e.g., “MPI version 2.0” refers to the standard in effect at the time of writing this document, while “current version of MPI” refers to further versions that take effect during the lifetime of this subcontract.

## TLCC2 Statement of Work

<b>EDAC</b>	Error Detection and Correction (EDAC) software based on the BlueSmoke technology ( <a href="http://www.sourceforge.net/projects/bluesmoke/">http://www.sourceforge.net/projects/bluesmoke/</a> )
<b>Fully supported</b> (as applied to system software and tools)	A product-quality implementation, documented and maintained by the HPC machine supplier or an affiliated software supplier.
<b>Gang Scheduling</b>	When a user job is scheduled to run, the Gang scheduler must contemporaneously allocate to CPUs all the threads and processes within that job (either within an SMP or within the cluster of SMPs). This scheduling capability must control all threads and processes within the SMP cluster environment.
<b>GFS</b>	Government Furnished Software (GFS) is software supplied to the Offeror by LLNS when TLCC2 build or installation takes place.
<b>Job</b>	A job is a cluster wide abstraction similar to a POSIX session, with certain characteristics and attributes. Commands will be available to manipulate a job as a single entity (including kill, modify, query characteristics, and query state). The characteristics and attributes required for each session type are as follows: 1) interactive session: an interactive session will include all cluster wide processes executed as a child (whether direct or indirect through other processes) of a login shell and will include the login shell process as well. Normally, the login shell process will exist in a process chain as follows: init, inetd, [sshd   telnetd   rlogind   xterm   cron], then shell. 2) batch session: a batch session will include all cluster wide processes executed as a child (whether direct or indirect through other processes) of a shell process executed as a child process of a batch system shepherd process, and will include the batch system shepherd process as well. 3) ftp session: an ftp session will include an ftpd and all its child processes. 4) kernel session: all processes with a pid of 0. 5) idle session: this session does not necessarily actually consist of identifiable processes. It is a pseudo-session used to report the lack of use of resources. 6) system session: all processes owned by root that are not a part of any other session.
<b>Lustre</b>	Lustre is an open source cluster wide file system based on object technology. See <a href="http://www.lustre.org">www.lustre.org</a> for more details. SU delivered to LLNL and Sandia will be configured with hardware and software to provide a Lustre cluster wide file system.
<b>MPI</b>	Message Passing Interface Version 1.2 or later. See, for example, <a href="http://www-unix.mcs.anl.gov/mpi/mpich/">http://www-unix.mcs.anl.gov/mpi/mpich/</a> , or <a href="http://www.mpi-forum.org/docs/mpi-20-html/mpi2-report.html">http://www.mpi-forum.org/docs/mpi-20-html/mpi2-report.html</a>
<b>OSPF</b>	Open Shortest Path First protocol. See, for example, <a href="http://www.ietf.org/rfc/rfc2328.txt">http://www.ietf.org/rfc/rfc2328.txt</a>
<b>Panasas PanFS</b>	PanFS is a proprietary cluster wide file system hardware and software solution based on industry standard object and interface specifications. See <a href="http://www.panasas.com">www.panasas.com</a> for more details. SU delivered to LANL will be configured with hardware and software to provide a Panasas cluster wide file system.

## TLCC2 Statement of Work

<b>Published</b> (as applied to APIs):	Where an API is not required to be consistent across platforms, the capability lists it as “published,” referring to the fact that it will be documented and supported, although it will be Offeror- or even platform-specific.
<b>Single-point control</b> (as applied to tool interfaces)	Refers to the ability to control or acquire information on all processes/PEs using a single command or operation.
<b>SNMP</b>	Simple Network Management Protocol
<b>Standard</b> (as applied to APIs)	Where an API is required to be consistent across platforms, the reference standard is named as part of the capability. The implementation will include all routines defined by that standard (even if some simply result in no-ops on a given platform).
<b>SWL</b>	Synthetic WorkLoad (SWL) is a set of applications representative of Tri-Laboratory workload used with Gazebo test harness to stress test the SU and clusters of SU aggregations. This SWL will only contain unclassified codes that are not export controlled.
<b>XXX-compatible</b> (as applied to system software and tool definitions)	Requires that a capability be compatible, at the interface level, with the referenced standard, although the lower-level implementation details will differ substantially (e.g., “NFSv4-compatible” means that the distributed file system will be capable of handling standard NFSv4 requests, but need not conform to NFSv4 implementation specifics).

**End of Section 7**